arXiv:1411.1124v2 [cs.DS] 7 Aug 2016

Nearly-Linear Time Packing and Covering LP Solvers

(Achieving Width-Independence and $O(1/\varepsilon)$-Convergence)

Zeyuan Allen-Zhu
zeyuan@csail.mit.edu
Princeton University

Lorenzo Orecchia
orecchia@bu.edu
Boston University

November 4, 2014*

Abstract

Packing and covering linear programs (PC-LPs) form an important class of linear programs (LPs) across computer science, operations research, and optimization. In 1993, Luby and Nisan [22] constructed an iterative first-order algorithm for approximately solving PC-LPs in *nearly-linear time*, independently of the program's *width*, i.e., the largest value in the program's description. Unfortunately, all existing nearly-linear-time algorithms [4, 9, 21, 29, 33, 34] for solving PC-LPs require time at least proportional to ε^{-2} to produce a $1 \pm \varepsilon$ approximation to the optimum.

In this paper, we break this longstanding barrier by designing a packing solver that runs in time $\widetilde{O}(N\varepsilon^{-1})$ and covering LP solver that runs in time $\widetilde{O}(N\varepsilon^{-1.5})$. Our packing solver can be extended to run in time $\widetilde{O}(N\varepsilon^{-1})$ for a class of well-behaved covering programs. In a follow-up work, Wang *et al.* [32] showed that all covering LPs can be converted into well-behaved ones by a reduction that blows up the problem size only logarithmically.

At high level, these two algorithms can be described as linear couplings of several first-order descent steps. This is the first application of our linear coupling technique (see [2]) to problems that are not amenable to blackbox applications known iterative algorithms in convex optimization.

*First version of this paper appeared on arXiv on this date. A 6-paged abstract of this paper has been presented at the STOC 2015 conference in Portland, OR. [3]

1 Introduction

A generic packing linear program (LP) takes the form $\max\{c^T x : Ax \leq b\}$ where $c \in \mathbb{R}^n_{\geq 0}$, $b \in \mathbb{R}^m_{\geq 0}$, and $A \in \mathbb{R}^{m \times n}_{\geq 0}$; similarly, a generic covering LP can be written as $\min\{b^T y : A^T y \geq c\}$, with the same requirements on A, b, and c. We denote by N the number of non-zero elements in matrix A.

We assume without loss of generality that the two LP programs are in their *standard forms*:

$$\text{Packing LP:} \qquad \max_{x \in \mathbb{R}^n_{\geq 0}} \{\mathbb{1}^T x : Ax \leq \mathbb{1}\} \ , \tag{1.1}$$

$$\text{Covering LP:} \qquad \min_{y \in \mathbb{R}^m_{\geq 0}} \{\mathbb{1}^T y : A^T y \geq \mathbb{1}\} \ . \tag{1.2}$$

Since the two programs are dual to each other, we denote by $\mathsf{OPT} \geq 0$ their shared optimal value. We say that x is a $(1-\varepsilon)$-approximation for the packing LP if $Ax \leq \mathbb{1}$ and $\mathbb{1}^T x \geq (1-\varepsilon)\mathsf{OPT}$, and x a $(1+\varepsilon)$-approximation for the covering LP if $A^T y \geq \mathbb{1}$ and $\mathbb{1}^T y \leq (1+\varepsilon)\mathsf{OPT}$.

In this paper, we study first-order iterative methods for solving packing and covering linear programs (PC-LPs) efficiently.[1] Of course, it is possible to adopt the Interior Point or Ellipsoid methods to obtain approximate solvers with a $\log(1/\varepsilon)$ dependence on the number of iterations. However, the computational cost of such algorithms is typically very high, as each iteration requires solving a linear system, and thus is not suitable for large-scale applications.

To address this issue, researchers have developed *iterative approximate* PC-LP solvers that achieve a better dependence on the problem size (e.g., nearly linear in N) at the cost of having a $\mathsf{poly}(1/\varepsilon)$ dependence on the approximation parameter ε. Such iterative solvers have been widely applied in approximation algorithms (e.g., MinSetCover [22], MaxSet, MaxDiCut, Max-k-CSP [30], bipartite matching), probabilistic checkable proofs [30], zero-sum matrix games [27], scheduling [29], graph embedding [29], flow controls [9, 10], auction mechanisms [35], wireless sensor networks [12], and many other areas. In addition, techniques developed in this line of research have also inspired important results on other fundamental algorithmic problems, such as the design of fast algorithms for multi-commodity flow problems [8, 17, 18, 23, 29] and the equivalence between QIP and PSPACE [19].

Previous approximate iterative solvers can be divided into two classes, *width-dependent* and *width-independent* solvers (see also Table 1).

Width-Dependent Solvers. These algorithms[2] usually require a running time that is at least N multiplied with $\rho\mathsf{OPT} \in [1, \infty)$, where ρ is the *width* of the program, i.e., the largest entry of matrix A. This results into a running time that is not polynomial, but only pseudo-polynomial.

More precisely, PC-LPs can be solved in $O(\frac{N\rho^2\mathsf{OPT}^2 \log m}{\varepsilon^2})$-time [29], or $O(\frac{N\rho\mathsf{OPT} \log m}{\varepsilon^2})$-time using negative-width techniques [6]. These algorithms strongly rely on multiplicative weight updates and only require "oracle-access" to the matrix A. When A is given explicitly like in this paper, the running time can be reduced to $O(\frac{N\rho\mathsf{OPT} \log m}{\varepsilon})$ by deploying more advanced optimization tools such as Nesterov's accelerated gradient method [27], or Nemirovski's mirror prox method [24].

Bienstock and Iyengar [11] have converted this dependence on $\rho\mathsf{OPT}$ into a more benign, yet linear dependence on n. More specifically, their running time is $O(\varepsilon^{-1}N\sqrt{Kn \log m})$ where K is the maximum number of non-zeros per row of A. This is $O(\varepsilon^{-1}Nn\sqrt{\log m})$ in the worst case. The

[1] Luby and Nisan, who originally studied iterative solvers for this class of problems [22], dubbed them *positive LPs*. However, the class of LPs with non-negative constraint matrices is slightly larger, including mixed-packing-and-covering LPs. For this reason, we prefer to stick to the PC-LP terminology.

[2] Most width-dependent solvers are studied under the minmax form of PC-LPs, which is $\min_{x \geq 0,\ \mathbb{1}^T x = 1} \ \max_{y \geq 0,\ \mathbb{1}^T y = 1} \ y^T A x$, whose optimal value equals $1/\mathsf{OPT}$. Their approximation guarantees are often written in terms of *additive* error. We have translated their performances to multiplicative error for a fair comparison.

Paper	Running Time	Width Independent?
[29]	$O(N \times \frac{\rho^2 \mathsf{OPT}^2 \log m}{\varepsilon^2})$	no
[6]	$O(N \times \frac{\rho \mathsf{OPT} \log m}{\varepsilon^2})$	no
[24, 27]	$O(N \times \frac{\rho \mathsf{OPT} \log m}{\varepsilon})$	no
[11]	$O(N \times \frac{\sqrt{Kn} \log m}{\varepsilon})$	no
[14, 25]: packing LP	$\widetilde{O}(N \times (n + \frac{\sqrt{n}}{\varepsilon}))$	no
[4, 7, 9, 10, 22, 32, 33]	$O(N \times \frac{\log^2 N \log(1/\varepsilon)}{\varepsilon^2})$ at best	yes
[33]	$O((md + N) \times \frac{\log N}{\varepsilon^2})$ [3]	yes
[9, 10]	$O(nm \times \frac{\log N}{\varepsilon^2})$	yes
[34]	$O(N \times \frac{\log N}{\varepsilon^2})$	yes
[21]	$O(N + (n+m) \times \frac{\log N}{\varepsilon^2})$	yes
our Theorem 3.2 packing LP	$O(N \times \frac{\log N \log \varepsilon^{-1}}{\varepsilon})$	yes
our Theorem 5.3 well-behaved covering LP	$O(N \times \frac{\log N \log \varepsilon^{-1}}{\varepsilon})$	yes
our Theorem 6.4 covering LP	$O(N \times \frac{\log N \log \varepsilon^{-1}}{\varepsilon^{1.5}})$	yes

Table 1: Comparisons among iterative approximate solvers for packing and covering LPs. The width $\rho \in [1/\mathsf{OPT}, \infty)$ is defined as the largest entry of the constraint matrix A.

results of [14, 25] have improved this convergence rate (for packing LP only) to $\widetilde{O}(\varepsilon^{-1} N \sqrt{n})$, but at a cost of enduring an $\widetilde{O}(Nn)$-time preprocessing stage.

Width-Independent Solvers. In this paper, we are interested in a second, more efficient class of methods, i.e. *width-independent* solvers.[4]

This line of research was initiated by a seminal paper of Luby and Nisan [22], who gave an algorithm running in $O\big(\frac{N \log^2 N}{\varepsilon^4}\big)$ time with no dependance on the width ρ. This is also the first *nearly-linear-time* approximate solver for PC-LPs, and also the first to run in parallel in nearly-linear-work and polylogarithmic depth.

The parallel algorithm of Luby and Nisan was extended by a sequence of works [4, 7, 9, 31, 33]. Most notably, the algorithm of Wang *et al.* [31] runs in $O(\frac{\log^2 N \log(1/\varepsilon)}{\varepsilon^2})$ iterations, each costing a matrix-vector multiplication that can be implemented in $O(N)$ total work and logarithmic depth.

The ideas of Luby and Nisan also led to sequential width-independent, nearly-linear time PC-LP solvers [9, 10, 21, 33, 34]. Most notably, the algorithm of Koufogiannakis and Young [21] runs in time $O\big(N + \frac{\log N}{\varepsilon^2} \times (n+m)\big)$.

Despite the amount of work in this area, the $O(1/\varepsilon^2)$ convergence rate was established in 1997 [9, 10] and has not been improved since then. On a separate note, Klein and Young [20] showed that all Dantzig-Wolfe type algorithms have to suffer from a $O(1/\varepsilon^2)$ convergence rate. This lack of progress constitutes a significant limitation, as the ε^{-2}-dependence (also known as the $1/\sqrt{T}$ convergence) on the approximation parameter ε is particularly poor.

[3] The parameter d is the maximum number of constraints each variable is in; md may be larger than N.

[4] Some of these solvers still have a polylog(ρ) dependence. Since each occurrence of $\log(\rho)$ can typically be replaced with $\log(nm)$ after slightly modifying the matrix A, we have done so in Table 1 for a fair comparisons.

1.1 Our Results

Packing LP. We present an algorithm `PacLPSolver` that can be implemented to run in $O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon}N)$ total time. This gives the first width-independent, and the first nearly-linear time solver for packing LP with an ε^{-1} convergence (i.e., an $1/T$ convergence). In contrast, no width-independent algorithm has achieved any convergence rate faster than ε^{-2} (or equivalently $O(1/\sqrt{T})$) before our work.

Interestingly, the maximum (weighted) bipartite matching is just one instance of a packing LP. As a consequence, our `PacLPSolver` algorithm finds an approximate maximum bipartite matching in time $\widetilde{O}(m\varepsilon^{-1})$ and an exact maximum matching in time $\widetilde{O}(m\sqrt{n})$.[5] Hence, this new matching algorithm, which arises purely from convex-optimization arguments, matches the running time of the best known combinatorial algorithm for maximum weighted bipartite matching (which is based on dynamic trees and data structures). Any further improvement over the dependence on ε^{-1} would result in a maximum matching algorithm that runs in time $m \cdot \widetilde{o}(\sqrt{n})$, which may require significantly different ideas.

Covering LP. A completely symmetric analysis of `PacLPSolver` gives rise to an algorithm $\texttt{CovLPSolver}^{\mathsf{wb}}$ with the same running time $O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon}N)$, but only solving *well-behaved* covering LP instances. At a high level, we say a covering LP instance is well-behaved if the constraint $A^T y \geq \mathbb{1}$ is "never redundant" (see Section 5 for the full definition). In particular, if the optimal solution y^* satisfies $C \cdot \mathbb{1} \geq A^T y^* \geq \mathbb{1}$ for some constant $C > 1$, then the covering LP is well-behaved.

As for the *general* covering LP *without* well-behavior assumptions, we propose a different algorithm `CovLPSolver` that runs in time $O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon^{1.5}}N)$. Again, we emphasize that no width-independent covering LP solver can achieve a convergence rate faster than ε^{-2} (or equivalently $O(1/\sqrt{T})$) before our work.

Remark. After the first version of this paper appeared on arXiv, Wang, Rao and Mahoney [32] showed that all covering LPs can be converted into well-behaved ones by a reduction that blows up the problem size only logarithmically. In other words, they obtained the ε^{-1} convergence rate through a reduction to $\texttt{CovLPSolver}^{\mathsf{wb}}$. Nevertheless, we point out that `CovLPSolver`, being a direct method for solving covering LPs, may be both of practical and theoretical interest.

1.2 Main Challenge and Our Approach

Width-Independence vs. Acceleration. Previous solvers for PC-LPs are based on standard iterative techniques in non-smooth optimization. These algorithms implicitly or explicitly smoothen the non-smooth objective, often by the entropy regularizer. The resulting convex objective is then minimized either via variations of full-gradient methods, yielding parallel algorithms, or via variations of coordinate-gradient methods, yielding sequential algorithms. The *main challenge* in previous work is to show that the width dependence can sometimes be removed for PC-LPs if the underlying minimization method is designed cleverly.

Note that, the slower the convergence rate is, the easier it becomes to design width-independent solvers. The $1/\varepsilon^4$-convergence solver of Awerbuch and Khandekar [7] and the $1/\varepsilon^3$-convergence solver of [4] are arguably the simplest width-independent solvers at this point.

In this paper, we achieve the $\frac{1}{\varepsilon}$ convergence that is typical for accelerated gradient descent over smoothened objectives [27], but without paying a possibly large width factor. The challenge in this

[5] It is not hard to turn an $\widetilde{O}(m\varepsilon^{-1})$ approximate algorithm into an $\widetilde{O}(m\sqrt{n})$ algorithm, see for instance [15].

approach is to preserve the width-independent behavior and the accelerated rate *at the same time*. We stress here that our algorithm is *not an instance* of any known variant of accelerated gradient descent[6]. Moreover, the incorporation of width-independence and Nesterov's acceleration requires significant effort, as witnessed by the lack of progress on this problem for the last 15 years.

Our result also overcomes the $1/\varepsilon^2$ lower bound of Klein and Young [20] since the underlying primal-dual approach of accelerated gradient descent falls outside of the class of Dantzig-Wolfe methods.

Our High-Level Approach. Our approach is based on an improved formalization of the properties of the smoothened PC-LP objective, together with our linear-coupling framework for designing efficient first-order methods [2].

The improved formalization shows that our smoothened objective satisfies either the classical condition for Lipschitz smoothness or a different condition based on multiplicative change.. This formalization greatly clarifies why width-independent algorithms exist in the first place. See Lemma 2.6 and the related discussion for more details.

We introduced the linear-coupling framework in previous work [2] providing a different interpretation of Nesterov's acceleration for smooth optimization [26]. In a nutshell, the linear-coupling framework allows us to construct accelerated algorithms by coupling the executions of a gradient descent algorithm, yielding iterates $\{\mathsf{y}_k\}$ and a mirror descent step algorithm, with iterates $\{\mathsf{z}_k\}$. The name "linear coupling" stems from the fact that, at iteration $k+1$, the gradient of the objective is queried at a point x_{k+1}, which is a linear combination of gradient and mirror steps, i.e., $\mathsf{x}_{k+1} = (1-\tau)\cdot z_k + (1-\tau)\cdot y_k$.

In this paper, we apply the linear-coupling framework in a very non-trivial manner. We design a new gradient and a new mirror descent step, each of them very specific to the underlying PC-LP problem. We also perform a coupling step $\mathsf{x}_{k+1} = (1-\tau)\cdot z_k + (1-\tau)\cdot y_k$, but we need to design a very different analysis (see Section 3.4) to preserve width independence. None of these components has appeared in our prior work [2].

1.3 Overview of Our Algorithms

Our Packing LP Solver. To solve packing LP, we minimize a relaxed version of the original LP, where the hard constraint $Ax \leq 1$ is regularized by entropy and replaced by an exponential penalty function. In other words, we reduce the problem of approximately solving packing LP into approximately minimizing some function $f_\mu(x)$ over the positive orthant $x \geq 0$ (see (2.3)).[7] We emphasize that traditional accelerated gradient methods [26, 27] should not be applied directly to minimize f_μ: this is because the smoothness of f_μ depends on the width of the LP and thus the resulting algorithm shall become width-dependent and cannot run in nearly-linear time. We thus need fundamentally new ideas.

Our proposed `PacLPSolver` algorithm analogous to randomized coordinate descent (cf. [5, 28]). At any point $x \geq 0$ of an iteration, it randomly picks a coordinate $i \in [n]$ and computes the coordinate gradient $\nabla_i f(x) \in [-1, \infty)$. A crucial step in `PacLPSolver` and its analysis is the decomposition of the gradient $\nabla_i f(x)$ into a large component $\eta \in [0, \infty)$ and a small component $\xi \in [-1, 1]$, satisfying $\nabla_i f(x) = \eta + \xi$. Then, `PacLPSolver` focuses only ξ, which we refer to as the

[6]This can be verified by observing that our already-smoothened objective $f_\mu(x)$ is not globally Lipschitz smooth, so that one cannot apply accelerated gradient descent directly.

[7]The explcit use of regularization to smoothen the PC-LP objective has been introduced by the same authors in previous work on parallel solvers [4]. However, the techniques in that paper only lead to very slow sequential solvers (see Table 1).

truncated gradient. This gradient-truncation technique was first developed in our prior work [4], but for a different purpose (to ensure parallelism) and not applied to coordinate gradient.

The rational behind gradient truncation is the following. If all the gradients were small (and thus truncation were not necessary), existing results on accelerated gradient descent, such as our linear coupling interpretation [2], would already yield the desired $(1/\varepsilon)$-convergence. The presence of large gradients η (and hence the possibility of rapidly changing gradients) makes the function f_μ non-smooth. To deal with this issue, we handle large gradients separately and only take steps with respect to the truncated gradient ξ.

More specifically, using the language of linear coupling developed in [2], we perform a gradient descent step and a mirror descent step both with respect to ξ. The gradient descent step helps us handle the large gradient η and thus we can assume as if η were 0. Next, we use linear coupling to combine gradient and mirror descent to take care of the truncated gradient ξ and provide an accelerated convergence rate $1/\varepsilon$.[8]

To push through the idea sketched above, we need to make some additional effort. Most notably, we impose a box constraint on x, by requiring each x_i to be upper bounded by a carefully chosen constant c_i. While this constraint $x_i \le c_i$ is provably redundant from the viewpoint of minimizing $f_\mu(x)$, it is necessary for (the mirror descent part of) linear coupling to work.

Our Covering LP Solver. Due to the randomness used in `PacLPSolver`, one cannot extract from it an (approximate) covering LP solution although packing and covering LPs are dual of each other. Therefore, we have to develop a direct solver for covering LP.

One way to build a direct solver is to turn `PacLPSolver` into its symmetric form and we call this algorithm $\texttt{CovLPSolver}^{\mathsf{wb}}$. Unfortunately, because the box constraint $x_i \in [0, c_i]$ is still needed in this symmetric version, the resulting algorithm only works when the covering LP instance is well-behaved, meaning that an optimal solution x^* must sit in this box constraint as well.

Another way to build it is to get rid of the box constraint. In a separate algorithm `CovLPSolver`, we replace the box with a (much larger) simplex which x^* unconditionally lies in. To take care of the extra search space (in the simplex minus the box), we have to change and "weaken" the underlying gradient and mirror descent steps as well. The final analysis again relies on linear coupling, but suffers from a factor $1/\sqrt{\varepsilon}$ loss in the convergence rate.

1.4 Roadmap

We relax the packing LP problem in Section 2, and provide our packing LP solver in Section 3. We relax the covering LP problem in Section 4, and provide our covering LP solver in the well-behaved case in Section 5. In Section 6, we provide our full covering LP solver.

2 Relaxation of the Packing Linear Program

Recall that, for input matrix $A \in \mathbb{R}^{m\times n}_{\ge 0}$, the packing LP in its standard form is $\max_{x\ge 0}\{\mathbb{1}^T x : Ax \le \mathbb{1}\}$. Let us denote by OPT the optimal value of this linear program, and x^* any optimal solution. We say that x is a $(1-\varepsilon)$-approximation for the packing LP if $Ax \le \mathbb{1}$ and $\mathbb{1}^T x \ge (1-\varepsilon)\mathsf{OPT}$.

[8] An alternative way to view this is to perform three steps in an iteration, two gradient descent and one mirror descent. One gradient descent takes care of the large gradient η, and the second gradient descent together with mirror descent take care of the truncated gradient ξ. We have simplified the algorithm so the two gradient steps become the same. This is the first time in the literature where 3 descent steps were coupled together to provide a fast convergence rate, as compared to 2 in accelerated gradient descent. In a separate paper of the first author, an accelerated method for stochastic gradient descent was obtained by explicitly performing three update steps and coupling them together.

Throughout this paper, we use the indices $i \in [n]$ to denote the columns of A, and the indices $j \in [m]$ to denote the rows of A. We let $A_{:i}$ be the i-th column vector of A, and $A_{j:}$ the j-th row vector of A. Given any vector x, we denote by $\|x\|_A = \sqrt{\sum_{i\in[n]} x_i^2 \cdot \|A_{:i}\|_\infty}$ the A-norm of x. By simple scaling, we can assume without loss of generality that[9]

$$\min_{i\in[n]}\{\|A_{:i}\|_\infty\} = 1 \ . \tag{2.1}$$

We restrict the domain of x and the range of OPT as follows.

Fact 2.1. *Define the bounding box* $\Delta_{\mathsf{box}} \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n : x_i \in \big[0, \frac{1}{\|A_{:i}\|_\infty}\big]\}$. *Under assumption (2.1), we have* $\mathsf{OPT} \in [1, n]$ *and* $\{x : x \geq 0 \wedge Ax \leq \mathbb{1}\} \subseteq \Delta_{\mathsf{box}}$.

Proof. Suppose that i^* is the column that achieves the smallest infinite norm $\|A_{:i}\|_\infty$ over all columns. Letting x be such that $x_i = 1$ at $i = i^*$ and $x_i = 0$ at $i \neq i^*$, we claim that x is a feasible solution for the packing LP (1.1), simply because $\|A_{:i^*}\|_\infty = 1$ according to (2.1). This feasible solution x yields an objective value $\mathbb{1}^T x = 1$, proving that $\mathsf{OPT} \geq 1$. On the other hand, for any solution $x \geq 0$ satisfying $Ax \leq \mathbb{1}$, we must have $x_i \leq \frac{1}{\|A_{:i}\|_\infty}$ for each i. Therefore, $\mathbb{1}^T x \leq \sum_i \frac{1}{\|A_{:i}\|_\infty} \leq n$, proving that $\mathsf{OPT} \leq n$.

The inclusion $\{x : x \geq 0 \wedge Ax \leq \mathbb{1}\} \subseteq \Delta_{\mathsf{box}}$ is obvious, since the constraints $x \geq 0$ and $Ax \leq \mathbb{1}$ together imply $x_i \leq \frac{1}{\|A_{:i}\|_\infty}$ for every $i \in [n]$. □

This bounding-box constraint allows us to focus only on searching x in Δ_{box}.

Our Regularized Objective. We now introduce the smoothed objective $f_\mu(x)$ that we minimize over Δ_{box} in order to approximately solve packing LP. At a high level, this objective $f_\mu(x)$ turns each row of the hard, non-smooth LP constraint $Ax \leq \mathbb{1}$ into an exponential penalty function so that we only need to require $x \in \Delta_{\mathsf{box}}$ throughout the algorithm.

Formally, the packing LP can be written as the following minimization problem by introducing the Lagrangian variable $y \in \mathbb{R}^m$:

$$\min_{x\in\Delta_{\mathsf{box}}} \big\{ -\mathbb{1}^T x \ + \ \max_{y\geq 0}\{y^T Ax - \mathbb{1}^T y\}\big\} \ . \tag{2.2}$$

The problem can be now smoothened by introducing a strongly concave regularizer over $y \geq 0$.

This is regularizer is usually taken to be the entropy function over all possible $y \geq 0$ satisfying $\mathbb{1}^T y = 1$, which yields the width-independent solvers in for instance [27] and [24], and is closely related to the so-called multiplicative weight update methods in [6].

In this paper, we instead take this regularizer to be the generalized entropy $H(y) = -\sum_{j=1}^m y_j \log y_j + y_j$ over the first orthant $y \geq 0$, and minimize the following smoothened objective $f_\mu(x)$ over $x \in \Delta_{\mathsf{box}}$:

$$f_\mu(x) \stackrel{\text{def}}{=} -\mathbb{1}^T x \ + \ \max_{y\geq 0}\{y^T Ax - \mathbb{1}^T y + \boxed{\mu \cdot H(y)}\} \ . \tag{2.3}$$

Above, $\mu > 0$ is some smoothing parameter to be chosen later. By explicitly computing the maximization over $y \geq 0$, $f_\mu(x)$ can be rewritten as

Lemma 2.2. $f_\mu(x) = \mu \sum_{j=1}^m e^{\frac{1}{\mu}((Ax)_j - 1)} - \mathbb{1}^T x$.

We study the *minimization* problem on $f_\mu(x)$ over $x \in \Delta_{\mathsf{box}}$. Intuitively $f_\mu(x)$ captures the original packing LP (1.1) as follows. Firstly, since we want to maximize $\mathbb{1}^T x$, the negative term $-\mathbb{1}^T x$ shows up in $f_\mu(x)$. Secondly, if a packing constraint $j \in [m]$ is violated by ε, that is,

[9] If $\min_{i\in[n]}\{\|A_{:i}\|_\infty\} = 0$ then the packing LP is unbounded so we are done. Otherwise, if $\min_{i\in[n]}\{\|A_{:i}\|_\infty\} = v > 0$ we scale all entries of A by $1/v$, and scale OPT by v.

$(Ax)_j \geq 1+\varepsilon$, the exponential penalty in $f_\mu(x)$ introduces a penalty at least $e^{\varepsilon/\mu}$; this will be a large penalty if $\mu \leq O(\varepsilon/\log n)$. Notice that this smoothed objective also appeared in previous works [4], albeit without this smoothening interpretation and without the constraint $x \in \Delta_{\mathsf{box}}$.

This $f_\mu(x)$ incurs a regularization error because we are now solving an objective different from the original LP. The next proposition bounds this regularization error following a similar treatment in a previous paper of the authors [4] and is proved in Appendix A for completeness' sake.

Proposition 2.3. *Let $\mu = \frac{\varepsilon}{4\log(nm/\varepsilon)}$ and x^* be an optimal solution for the packing LP (1.1). Then:*

(a) $f_\mu(u^*) \leq -(1-\varepsilon)\mathsf{OPT}$ for $u^* \stackrel{\text{def}}{=} (1-\varepsilon/2)x^* \in \Delta_{\mathsf{box}}$.
(b) $f_\mu(x) \geq -(1+\varepsilon)\mathsf{OPT}$ for every $x \in \Delta_{\mathsf{box}}$.
(c) If $x \in \Delta_{\mathsf{box}}$ satisfies $f_\mu(x) \leq -(1-O(\varepsilon))\mathsf{OPT}$, then $\frac{1}{1+\varepsilon}x$ is a $(1-O(\varepsilon))$-approximate solution to the packing LP.

Remark 2.4. Our box constraint $x \in \Delta_{\mathsf{box}}$ is almost redundant from the viewpoint of minimizing $f_\mu(x)$: whenever $x \geq 0$ and $f_\mu(x) \leq 0$, one should automatically have $x_i \leq \frac{1+\varepsilon}{\|A_{:i}\|_\infty}$. However, this redundant constraint is crucial in maintaining the mirror-descent component of our algorithm bounded.

Smoothness Properties. We now show $f_\mu(x)$ enjoys a number of non-standard smoothness properties.

It is first easy to write the gradient and Hessian of $f_\mu(x)$ in the following closed forms:

Fact 2.5. $\nabla f_\mu(x) = A^T p(x) - \mathbb{1}$ *and* $\nabla^2 f_\mu(x) = \frac{1}{\mu}A^T \mathrm{diag}\{p(x)\}A$,
where $p_j(x) \stackrel{\text{def}}{=} e^{\frac{1}{\mu}((Ax)_j - 1)}$.

By staring at these close forms, we note that $f_\mu(x)$ is not Lipschitz-smooth: for instance, each $\nabla^2_{ii} f_\mu(x)$ can go to infinity so the spectral norm of $\nabla^2 f_\mu(x)$ is unbounded.

However, the non-negativity of A guarantees that whenever $\nabla^2_{ii} f_\mu(x)$ is large for some coordinate i, the corresponding entry of the gradient $\nabla_i f_\mu(x)$ must also be large. In other words, if the gradient of $f_\mu(x)$ is poorly smooth along coordinate i at x, the gradient along i at x must be very large. This still allows us to take a larger step in direction $\mathbf{e}_i$ than traditionally allowed by coordinate descent.

The above intuition is formalized in the next lemma. The first half of the lemma is the same as the traditional coordinate Lipschitz-smoothness property, but holds only conditionally and not for all $x \in \Delta_{\mathsf{box}}$. The second half of the lemma is a salient characteristic of this work and requires the non-negativity of A.

Lemma 2.6. *Define the smoothness parameter $L \stackrel{\text{def}}{=} \frac{4}{\mu}$. Then, for every $x \geq 0$, every $i \in [n]$, and every $\lambda \leq \frac{1}{L\|A_{:i}\|_\infty}$:*

(a) If $|\nabla_i f_\mu(x)| \leq 1$, then $\left|\nabla_i f_\mu(x+\lambda\mathbf{e}_i) - \nabla_i f_\mu(x)\right| \leq L\|A_{:i}\|_\infty \cdot |\lambda|$.
(b) If $\nabla_i f_\mu(x) \geq 1$, then $\nabla_i f_\mu(x+\lambda\mathbf{e}_i) \geq \left(1 - \frac{\|A_{:i}\|_\infty L}{2}|\lambda|\right)\nabla_i f_\mu(x)$.

These smoothness properties will be crucial in applying gradient descent arguments in Section 3.3, and are the main motivation for us to adopt the $\|\cdot\|_A$ norm for our proposed algorithms. Their proofs require manipulations of the Hessian and are included in Appendix A.

Initialization. Iterative methods require to choose a starting point, and we use the following one

Fact 2.7. *Defining $x_i^{\mathsf{start}} \stackrel{\text{def}}{=} \frac{1-\varepsilon/2}{n\|A_{:i}\|_\infty}$ for for each $i \in [n]$, we have $x^{\mathsf{start}} \in \Delta_{\mathsf{box}}$ and $f_\mu(x^{\mathsf{start}}) \leq -\frac{1-\varepsilon}{n}$.*

(The proof of Fact 2.7 is straightforward but anyways included in Appendix A.)

Algorithm 1 $\texttt{PacLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$

Input: $A \in \mathbb{R}^{m\times n}_{\geq 0}, x^{\mathsf{start}} \in \Delta_{\mathsf{box}}, \varepsilon \in (0, 1/10]$.
Output: $x \in \Delta_{\mathsf{box}}$.

$\mu \leftarrow \frac{\varepsilon}{4\log(nm/\varepsilon)}$, $L \leftarrow \frac{4}{\mu}$, $\tau \leftarrow \frac{1}{3\cdot nL}$ and $\alpha_0 \leftarrow \frac{1}{nL}$. ▷ parameters

$T \leftarrow \lceil 3nL\log(1/\varepsilon)\rceil = O(n \cdot \frac{\log(nm/\varepsilon)\cdot\log(1/\varepsilon)}{\varepsilon})$. ▷ number of iterations

$\mathsf{x}_0 = \mathsf{y}_0 \leftarrow x^{\mathsf{start}}$, $\mathsf{z}_0 \leftarrow 0$.

for $k \leftarrow 1$ **to** T **do**

  $\alpha_k \leftarrow \frac{1}{1-\tau}\,\alpha_{k-1}$

  $\mathsf{x}_k \leftarrow \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$.

  Randomly select $i \in [n]$ uniformly at random.

  Define vector $\xi_k^{(i)}$ to be all-zero except at coordinate i, where it equals $\mathbb{T}^{\mathsf{p}}(\nabla_i f_\mu(\mathsf{x}_k))$.

  $\mathsf{z}_k \leftarrow \mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z\in\Delta_{\mathsf{box}}} \left\{\frac{1}{2}\|z - \mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k\xi_k^{(i)}, z\rangle\right\}$. ▷ See Proposition 3.7

  $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k + \frac{1}{n\alpha_k L}(\mathsf{z}_k^{(i)} - \mathsf{z}_{k-1})$.

end for

return y_T.

3 Our Packing LP Solver

To describe our packing LP solver, we choose a thresholding function:

Definition 3.1. *The thresholding function* $\mathbb{T}^{\mathsf{p}}\colon [-1,\infty) \to [-1,1]$ *is*

$$\mathbb{T}^{\mathsf{p}}(v) \stackrel{\text{def}}{=} \begin{cases} v, & v \in [-1,1]; \\ 1, & v > 1. \end{cases}$$

Our algorithm $\texttt{PacLPSolver}$ starts with some initial vector $\mathsf{x}_0 = \mathsf{y}_0 = x^{\mathsf{start}}$ (introduced in Fact 2.7) and $\mathsf{z}_0 = 0$, and is divided into T iterations. In each iteration, we start by computing a weighted midpoint $\mathsf{x}_k \leftarrow \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$ for some parameter $\tau \in (0,1)$, and then proceed to compute y_k and z_k as follows.

- Select $i \in [n]$ uniformly at random, and let $\xi_k^{(i)} = (0,\ldots,0,\mathbb{T}^{\mathsf{p}}(v),0,\ldots,0)$ be the vector that is only non-zero at coordinate i, where $v = \nabla_i f_\mu(\mathsf{x}_k) = \sum_{j=1}^m A_{j,i} e^{\frac{1}{\mu}((A\mathsf{x}_k)_j - 1)} - 1 \in [-1,\infty)$.
- Perform a *mirror (descent) step* $\mathsf{z}_k \leftarrow \mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z\in\Delta_{\mathsf{box}}} \left\{\frac{1}{2}\|z - \mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k\xi_k^{(i)}, z\rangle\right\}$ for some parameter $\alpha_k \ll 1/n$ to be chosen later.
- Perform a *gradient (descent) step* $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k + \frac{1}{n\alpha_k L}(\mathsf{z}_k^{(i)} - \mathsf{z}_{k-1})$.

Above, the reason that the the two steps on y_k and z_k are named after "gradient step" and "mirror step" will become clear in the follow-up subsections. We use the superscript $^{(i)}$ on $\xi_k^{(i)}$, $\mathsf{y}_k^{(i)}$ and $\mathsf{z}_k^{(i)}$ to emphasize that the value depends on the choice of i. We use generic parameters τ, α_k, T in the above description and their precise values are presented in Algorithm 1. [10]

3.1 Convergence Statement

In this section, we focus on proving the following main theorem in three steps.

Theorem 3.2. $\texttt{PacLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs some* y_T *satisfying*

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \leq -(1-3\varepsilon)\mathsf{OPT} \ .$$

[10] We encourage the readers to ignore their specific values for now. Our specific choices of the parameters shall become clear and natural at the end of this section, and be discussed whenever they are used.

At the same time, the running time is due to the following simple lemma:

Lemma 3.3. *Each iteration of* `PacLPSolver` *can be implemented to run in expected $O(N/n)$ time. The total expected running time is $O(TN/n)$.*

(The proof of Lemma 3.3 is not hard and included in Appendix F. The key idea used in the implementation is to compute x_k and y_k only *implicitly*. For instance, explicitly maintaining x_k and computing $p(x_k)$ require $O(N)$ time per iteration, but representing x_k implicitly as a linear combination of two less-frequently-modified vectors reduces it to $O(N/n)$.)

It is straightforward (see Appendix B) to use Markov inequality to turn the expected guarantee in Theorem 3.2 into a probabilistic one:

Corollary 3.4. *With probability at least 9/10,* $\texttt{PacLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs a $(1-O(\varepsilon))$ approximate solution to the packing LP program. The expected running time is $O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon}N)$.*

We also note that our iterates $\mathsf{x}_k, \mathsf{y}_k, z_k$ never leave the bounding box Δ_{box}:

Lemma 3.5. *We have $\mathsf{x}_k, \mathsf{y}_k, \mathsf{z}_k \in \Delta_{\mathsf{box}}$ for all $k = 0, 1, \ldots, T$.*

(The proof of Lemma 3.5 is included in Appendix B, and the main idea behind this lemma is related to randomized coordinate descent [16].)

We are now ready to prove Theorem 3.2 in the subsequent subsections.

3.2 Step 1: Mirror Descent Guarantee

Since our update $\mathsf{z}_k^{(i)} = \arg\min_{z\in\Delta_{\mathsf{box}}} \left\{\frac{1}{2}\|z-\mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k\xi_k^{(i)}, z\rangle\right\}$ —see Line 9 of `PacLPSolver`— is written in the form of a *mirror descent step*, the following inequality is a classical upper bound on the "regret" of mirror descent. Its proof can be found in Appendix B.

Lemma 3.6 (mirror descent)**.** *For every $u \in \Delta_{\mathsf{box}}$, it satisfies $\left\langle n\alpha_k\xi_k^{(i)}, \mathsf{z}_{k-1}-u\right\rangle \le n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)}\right\rangle + \frac{1}{2}\|\mathsf{z}_{k-1}-u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)}-u\|_A^2$.*

Although defined in a variational way, it is perhaps beneficial to explicitly describe how to implement this mirror step. Its proof is straightforward but included in Appendix B.

Proposition 3.7. *If $\mathsf{z}_{k-1} \in \Delta_{\mathsf{box}}$, the minimizer $z = \arg\min_{z\in\Delta_{\mathsf{box}}}\left\{\frac{1}{2}\|z-\mathsf{z}_{k-1}\|_A^2 + \langle\delta\mathbf{e}_i, z\rangle\right\}$ for any scalar $\delta\in\mathbb{R}$ and basis vector $\mathbf{e}_i$ can be computed as follows:*

1. *$z \leftarrow \mathsf{z}_{k-1}$.*
2. *$z_i \leftarrow z_i - \delta/\|A_{:i}\|_\infty$.*
3. *If $z_i < 0$, then $z_i \leftarrow 0$; if $z_i > 1/\|A_{:i}\|_\infty$, $z_i \leftarrow 1/\|A_{:i}\|_\infty$.*
4. *Return z.*

As a simple corollary, we have the following fact

Fact 3.8. $|\mathsf{z}_{k,i}^{(i)} - \mathsf{z}_{k-1,i}| \le \frac{n\alpha_k|\xi_{k,i}^{(i)}|}{\|A_{:i}\|_\infty}$ *and* $|\mathsf{y}_{k,i}^{(i)} - \mathsf{x}_{k,i}| = \frac{1}{n\alpha_k L}|\mathsf{z}_{k,i}^{(i)} - \mathsf{z}_{k-1,i}| \le \frac{|\xi_{k,i}^{(k)}|}{L\|A_{:i}\|_\infty} \le \frac{1}{L\|A_{:i}\|_\infty}$.

3.3 Step 2: Gradient Descent Guarantee

We call our update $\mathsf{y}_k^{(i)} \leftarrow \mathsf{x}_k + \frac{1}{n\alpha_k L}(\mathsf{z}_k^{(i)} - \mathsf{z}_{k-1})$ gradient step, because the following lemma guarantees $f_\mu(\mathsf{y}_k^{(i)}) \le f_\mu(\mathsf{x}_k)$, that is, the objective only decreases; moreover, the objective decreases at least by $\frac{1}{2}\langle\nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)}\rangle$.

Lemma 3.9 (gradient descent)**.** *We have* $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq \frac{1}{2}\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)}\rangle$*. In particular, this implies* $f_\mu(\mathsf{x}_k) \geq f_\mu(\mathsf{y}_k^{(i)})$ *because* $\nabla_i f_\mu(\mathsf{x}_k)$ *and* $\mathsf{x}_{k,i} - \mathsf{y}_{k,i}^{(i)}$ *have the same sign, while* $\mathsf{x}_{k,\ell} = \mathsf{y}_{k,\ell}^{(i)}$ *for* $\ell \neq i$*.*

This Lemma 3.9, which is characteristic of the PC-LP setting, is very strong in the following sense. Even though the iterative update only depends on the truncated gradient ξ_k^i in `PacLPSolver`, the progress we make is a function of the true gradient $\nabla_i f_\mu(\mathsf{x}_k)$, including the large component that was discarded. This is possible because the smoothness guarantee of Lemma 2.6 allows us to take a long coordinate step even though the function is not Lipschitz-smooth. We include the detailed proof of Lemma 3.9 in Appendix B.

3.4 Step 3: Putting All Together

We denote by $\eta_k^{(i)} \in \mathbb{R}^n_{\geq 0}$ the vector that is only non-zero at coordinate i, and satisfies $\eta_{k,i}^{(i)} = \nabla_i f_\mu(\mathsf{x}_k) - \xi_{k,i}^{(i)} \in [0, \infty)$. In other words, the full gradient

$$\nabla f_\mu(\mathsf{x}_k) = \mathbf{E}_i[n\nabla_i f_\mu(\mathsf{x}_k)] = \mathbf{E}_i[n\eta_k^{(i)} + n\xi_k^{(i)}]$$

can be (in expectation) decomposed into the a large but non-negative component $\eta_k^{(i)} \in [0,\infty)^n$ and a small component $\xi_k^{(i)} \in [-1,1]^n$. Recall that $\eta_k^{(i)}$ is the part of the gradient that was truncated, and did not contribute to the descent steps (see Line 9 of `PacLPSolver`). Next, for any $u \in \Delta_{\mathsf{box}}$, we can use a basic convexity argument and the mirror descent lemma to compute that

$$\begin{aligned}
&\alpha_k(f_\mu(\mathsf{x}_k) - f_\mu(u)) \leq \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u\right\rangle \\
&= \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{z}_{k-1}\right\rangle + \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - u\right\rangle \\
&= \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{z}_{k-1}\right\rangle + \mathbf{E}_i\left[\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle + \left\langle n\alpha_k \xi_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle\right] \\
&\overset{①}{=} \frac{(1-\tau)\alpha_k}{\tau}\left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{y}_{k-1} - \mathsf{x}_k\right\rangle + \mathbf{E}_i\left[\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle + \left\langle n\alpha_k \xi_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle\right] \qquad (3.1) \\
&\overset{②}{\leq} \frac{(1-\tau)\alpha_k}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) \\
&\qquad + \mathbf{E}_i\Big[\boxed{\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)}\right\rangle} + \frac{1}{2}\|\mathsf{z}_{k-1} - u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)} - u\|_A^2\Big] \qquad (3.2)
\end{aligned}$$

Above, ① is because $\mathsf{x}_k = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$, which implies that $\tau(\mathsf{x}_k - \mathsf{z}_{k-1}) = (1-\tau)(\mathsf{y}_{k-1} - \mathsf{x}_k)$. ② uses convexity and Lemma 3.6. This above computation is motivated by [2], and as we shall see below, it allows one to linearly couple gradient and mirror steps.

Intuitively, the first (non-negative) term in the box of (3.2) is the regret introduced by the large gradient $\eta_k^{(i)}$. This part was truncated so did not contribute to the mirror step. The second (non-negative) term in the box is the quadratic width-term introduced by mirror descent on the small gradient $\xi_k^{(i)}$.

Now comes an important observation. As shown by Lemma 3.10 below, the performance of the gradient step —that is, the objective decrease of $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})$— is at least proportional to the loss incurred in the box. Intuitively, this means that the progress in the gradient step is so large that it outweights not only the width-term from mirror descent (as is typical in accelerated gradient analyses) but also the regret term introduced by $\eta_k^{(i)}$.

Lemma 3.10 (gradient descent total guarantee)**.** *For every* $u \geq 0$*,*

$$\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)}\right\rangle \leq 3n\alpha_k L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \ .$$

Since the proof of the above lemma is a careful case analysis and several simple applications of Lemma 3.9, we defer it to Appendix B. However, we remark that to properly upper bound $\langle n\alpha_k\eta_k^{(i)}, \mathsf{z}_{k-1}-u\rangle$, one needs to have some good upper bound the coordinates of z_{k-1}. This is exactly the place we need our redundant-constraint technique, which guarantees that each $\mathsf{z}_{k-1,i} \le \frac{1}{\|A_{:i}\|_\infty}$.

Plugging the above lemma into (3.2), we have

$$\begin{aligned}
&\alpha_k(f_\mu(\mathsf{x}_k) - f_\mu(u)) \le \big\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u\big\rangle \\
&\overset{①}{\le} \frac{(1-\tau)\alpha_k}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \mathbf{E}_i\Big[3n\alpha_k L\cdot(f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) + \frac{1}{2}\|\mathsf{z}_{k-1}-u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k - u\|_A^2\Big] \\
&\overset{②}{\le} \alpha_k f_\mu(\mathsf{x}_k) + \big(3n\alpha_k L - \alpha_k\big) f_\mu(\mathsf{y}_{k-1}) + \mathbf{E}_i\Big[-3n\alpha_k L\cdot f_\mu(\mathsf{y}_k^{(i)}) + \frac{1}{2}\|\mathsf{z}_{k-1}-u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k - u\|_A^2\Big] \ .
\end{aligned} \tag{3.3}$$

Above, ① is because we have chosen α_k so that $n\alpha_k \le n\alpha_T = \frac{1}{\varepsilon L} \le \frac{1}{4}$; and ② is because we have chosen τ to satisfy $\frac{1}{\tau} = 3nL$.

Next, recall that we have picked α_k so that $(3nL-1)\alpha_k = 3nL\cdot\alpha_{k-1}$ in Algorithm 1. Telescoping (3.3) for $k = 1,\dots,T$ and choosing $u^* = (1-\varepsilon/2)x^*$, we have

$$-\sum_{k=1}^T \alpha_k f_\mu(u^*) \le 3f_\mu(\mathsf{y}_0) - 3n\alpha_T L\cdot\mathbf{E}[f_\mu(\mathsf{y}_T)] + \|\mathsf{z}_0 - u^*\|_A^2 \le -3n\alpha_T L\cdot\mathbf{E}[f_\mu(\mathsf{y}_T)] + \mathsf{OPT} \ .$$

Here, the second inequality is due to $f_\mu(\mathsf{y}_0) = f_\mu(x^{\mathsf{start}}) \le 0$ from Fact 2.7, and the fact that

$$\|\mathsf{z}_0 - u^*\|_A^2 = \|u^*\|_A^2 = \sum_{i=1}^n (u_i^*)^2\cdot\|A_{:i}\|_\infty \le \sum_{i=1}^n (x_i^*)^2\cdot\|A_{:i}\|_\infty \le \sum_{i=1}^n x_i^* = \mathsf{OPT} \ .$$

Finally, using the fact that $\sum_{k=1}^T \alpha_k = \alpha_T\cdot\sum_{k=0}^{T-1}\big(1-\frac{1}{3nL}\big)^k = 3n\alpha_T L\big(1-(1-\frac{1}{3nL})^T\big)$, we rearrange and obtain that

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \le \frac{\sum_k \alpha_k}{3n\alpha_T L} f_\mu(u^*) + \frac{1}{3n\alpha_T L}\mathsf{OPT} = \big(1-(1-\frac{1}{3nL})^T\big) f_\mu(u^*) + \frac{1}{3n\alpha_T L}\mathsf{OPT} \ .$$

Choosing $T = \lceil 3nL\log(1/\varepsilon)\rceil$ so that $\frac{1}{n\alpha_T L} = (1-\frac{1}{3nL})^T \le \varepsilon$. Combining this with the fact that $f_\mu(u^*) \le -(1-\varepsilon)\mathsf{OPT} < 0$ (see Proposition 2.3.a), we obtain

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \le (1-\varepsilon)f_\mu(u^*) + \varepsilon/3\cdot\mathsf{OPT} < -(1-3\varepsilon)\mathsf{OPT} \ .$$

Therefore, we have finished proving Theorem 3.2. □

4 Relaxation of the Covering Linear Program

As mentioned in the introduction, due to the randomized nature of `PacLPSolver`, we cannot directly infer from it a dual covering LP solution. Therefore, we have to work on a relaxed version of covering LP directly.

For input matrix $A \in \mathbb{R}_{\ge 0}^{m\times n}$, we rewrite the covering LP problem (1.2) as follows in order to be notationally similar to packing LP:

$$\min_{x\in\mathbb{R}_{\ge 0}^n}\{\mathbb{1}^T x \ : \ Ax \ge \mathbb{1}\} \ . \tag{4.1}$$

We denote by OPT the optimal value to this LP, and by x^* any of its optimal solutions. We say that x is a $(1+\varepsilon)$-approximation for the covering LP if $Ax \ge \mathbb{1}$ and $\mathbb{1}^T x \le (1+\varepsilon)\mathsf{OPT}$.

Again, we use indices $i \in [n]$ for the columns of A, and indices $j \in [m]$ for the rows of A. We denote by $A_{:i}$ the i-th column vector of A, and $A_{j:}$ the j-th row vector of A. We assume without

loss of generality by simple scaling that[11]

$$\min_{j\in[m]}\{\|A_{j:}\|_\infty\} = 1 \ . \tag{4.2}$$

In our covering LP solvers, we assume that an initial solution, achieving a constat approximation, is available to the algorithm. Such a solution can be obtained via the covering LP solver from Young [34] with constant ϵ in time $O(N \log N)$.

Definition 4.1. *Let $x^\sharp$ be a given 2-approximate solution to the covering problem given and let $\mathsf{OPT}' \stackrel{\text{def}}{=} \mathbb{1}^T x^\sharp \in [\mathsf{OPT}, 2\mathsf{OPT}]$.*

We now introduce the smoothed objective $f_\mu(x)$ that we are going to minimize in order to approximately solve the covering LP. This smoothed objective turns each row of the LP constraint $Ax \geq \mathbb{1}$ into an exponential penalty function. We skip the details regarding how it arises from a relaxation using the generalized entropy regularizer, because it is symmetric to Section 2.

Definition 4.2. *Letting parameter $\mu \stackrel{\text{def}}{=} \frac{\varepsilon}{4\log(nm/\varepsilon)}$, we define the smoothed objective $f_\mu(x)$ as*

$$f_\mu(x) \stackrel{\text{def}}{=} \mu \sum_{j=1}^{m} e^{\frac{1}{\mu}(1-(Ax)_j)} + \mathbb{1}^T x \ .$$

We begin with several simple properties about OPT and $f_\mu(x)$. They together imply that the minimum of $f_\mu(x)$ is around OPT, and if one approximately finds the minimum of $f_\mu(x)$ up to an additive error $O(\varepsilon\mathsf{OPT})$, this corresponds to a $(1+O(\varepsilon))$-approximate solution to the covering LP (4.1). The proofs are analogous to their counterparts in Section 2, and we include them in Appendix C for completeness' sake.

Proposition 4.3.

(a) $\mathsf{OPT} \in [1, m]$.
(b) $f_\mu(u^*) \leq (1+\varepsilon)\mathsf{OPT}$ *for* $u^* \stackrel{\text{def}}{=} (1+\varepsilon/2)x^*$.
(c) $f_\mu(x) \geq (1-\varepsilon)\mathsf{OPT}$ *for every* $x \geq 0$.
(d) For any $x \geq 0$ *satisfying* $f_\mu(x) \leq 2\mathsf{OPT}$, *we must have* $Ax \geq (1-\varepsilon)\mathbb{1}$.
(e) If $x \geq 0$ *satisfies* $f_\mu(x) \leq (1+O(\varepsilon))\mathsf{OPT}$, *then* $\frac{1}{1-\varepsilon}x$ *is a* $(1+O(\varepsilon))$*-approximate solution to the covering LP.*
(f) The gradient of $f_\mu(x)$ *can be written as*

$$\nabla f_\mu(x) = \mathbb{1} - A^T p(x) \quad \textit{where} \quad p_j(x) \stackrel{\text{def}}{=} e^{\frac{1}{\mu}(1-(Ax)_j)} \tag{4.3}$$

5 Our Covering LP Solver in the Well-Conditioned Case

Recall that in packing LPs, since it satisfies $0 \leq x_i^* \leq \frac{1}{\|A_{:i}\|_\infty}$ (see Fact 2.1), we can minimize f_μ over a bounding box Δ_{box}. Unfortunately, because we no longer have $x_i^* \leq \frac{1}{\|A_{:i}\|_\infty}$ in covering LPs, if one directly turns `PacLPSolver` into its symmetric covering LP solver, it is not sufficient to minimize $f_\mu(x)$ over Δ_{box}.

In this section, we show that this symmetric covering LP solver still solves all *well-behaved* covering LP instances. Specifically, we say that the covering LP problem is well-behaved if:[12]

[11] If $\min_{j\in[m]}\{\|A_{j:}\|_\infty\} = 0$ then the covering LP is infeasible so we are done. Otherwise, if $\min_{j\in[m]}\{\|A_{j:}\|_\infty\} = v > 0$ we scale all entries of A by $1/v$, and scale OPT by v.

[12] The constant 9 in this section can be replaced with any other constant greater than 1.

Algorithm 2 $\texttt{CovLPSolver}^{\mathsf{wb}}(A, x^{\mathsf{start}}, \varepsilon)$

Input: $A \in \mathbb{R}^{m\times n}_{\geq 0}, x^{\mathsf{start}} \in \Delta_{\mathsf{box}}, \varepsilon \in (0, 1/10]$.
Output: $x \in \Delta_{\mathsf{box}}$.

$\mu \leftarrow \frac{\varepsilon}{4\log(nm/\varepsilon)}$, $L \leftarrow \frac{4}{\mu}$, $\tau \leftarrow \frac{1}{21\cdot nL}$ and $\alpha_0 \leftarrow \frac{1}{nL}$. ▷ parameters
$T \leftarrow \lceil 21nL\log(1/\varepsilon)\rceil = O(n \cdot \frac{\log(nm/\varepsilon)\cdot\log(1/\varepsilon)}{\varepsilon})$. ▷ number of iterations
$\mathsf{x}_0 = \mathsf{y}_0 \leftarrow x^{\mathsf{start}}$, $\mathsf{z}_0 \leftarrow 0$.
for $k \leftarrow 1$ **to** T **do**
    $\alpha_k \leftarrow \frac{1}{1-\tau}\,\alpha_{k-1}$
    $\mathsf{x}_k \leftarrow \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$.
    Randomly select $i \in [n]$ uniformly at random.
    Define vector $\xi_k^{(i)}$ be all-zero except at coordinate i where it equals $-\mathbb{T}^{\mathsf{p}}(-\nabla_i f_\mu(\mathsf{x}_k))$.
    $\mathsf{z}_k \leftarrow \mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z\in\Delta_{\mathsf{box}}} \left\{\frac{1}{2}\|z - \mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k\xi_k^{(i)}, z\rangle\right\}$. ▷ See Proposition 3.7
    $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k + \frac{1}{n\alpha_k L}(\mathsf{z}_k^{(i)} - \mathsf{z}_{k-1})$.
end for
return y_T.

Assumption 5.1. *There exists some optimal covering LP solution x^* satisfying $x_i^* \leq \frac{9}{\|A_{:i}\|_\infty}$; and the starting point $x^\sharp$ satisfies $x_i^\sharp \leq \frac{9}{\|A_{:i}\|_\infty}$.*

For instance, well-behaved instances naturally arise from those where the constraints $Ax \geq \mathbb{1}$ are non-redundant. If the optimal covering LP solution x^* and the starting point $x^\sharp$ satisfy $\mathbb{1} \leq Ax^* \leq 9\cdot\mathbb{1}$ and $\mathbb{1} \leq Ax^\sharp \leq 9\cdot\mathbb{1}$, then Assumption 5.1 is satisfied.

Well-behaved covering LP problems immediately satisfy the following:

Fact 5.2. *Define $\Delta_{\mathsf{box}} \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n : x_i \in \left[0, \frac{10}{\|A_{:i}\|_\infty}\right]\}$. Under Assumption 5.1, we have $u^* \stackrel{\text{def}}{=} (1+\varepsilon/2)x^*$ and $x^{\mathsf{start}} \stackrel{\text{def}}{=} (1+\varepsilon/2)\cdot x^\sharp$ both reside in Δ_{box}. Also, it satisfies $f_\mu(x^{\mathsf{start}}) \leq 3\mathsf{OPT}$.*

Proof. We only prove $f_\mu(x^{\mathsf{start}}) \leq 3\mathsf{OPT}$ because the claims $u^*, x^{\mathsf{start}} \in \Delta_{\mathsf{box}}$ are trivial after noticing $\varepsilon \leq 1/10$. Using the fact that $Ax^{\mathsf{start}} - \mathbb{1} \geq (1+\varepsilon/2)Ax^\sharp - \mathbb{1} \geq \varepsilon/2\cdot\mathbb{1}$, we compute $f_\mu(x^{\mathsf{start}})$ as follows:

$$\begin{aligned} f_\mu(x^{\mathsf{start}}) &= \mu\sum_j e^{\frac{1}{\mu}(1-(Ax^{\mathsf{start}})_j)} + \mathbb{1}^T x^{\mathsf{start}} \leq \mu\sum_j e^{\frac{-\varepsilon/2}{\mu}} + 2\mathsf{OPT} \\ &\leq \frac{\mu m}{(nm)^2} + 2\mathsf{OPT} < 3\mathsf{OPT}\ . \end{aligned}$$

□

We now describe $\texttt{CovLPSolver}^{\mathsf{wb}}$ (which is a symmetric variant of `PacLPSolver`) that solves well-behaved covering LP problems, see Algorithm 2.

We start with the initial vector $\mathsf{x}_0 = \mathsf{y}_0 = x^{\mathsf{start}} = (1+\varepsilon/2)x^\sharp$ and $\mathsf{z}_0 = 0$. Then, $\texttt{CovLPSolver}^{\mathsf{wb}}$ is divided into T iterations. In each iteration, it computes a weighted midpoint $\mathsf{x}_k \leftarrow \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$ for some parameter $\tau \in (0,1)$, and then proceeds to compute y_k and z_k as follows.

- Select $i \in [n]$ uniformly at random, and let $\xi_k^{(i)} = (0,\ldots,0,-\mathbb{T}^{\mathsf{p}}(v),0,\ldots,0)$ be the vector that is only non-zero at coordinate i, where $v = -\nabla_i f_\mu(\mathsf{x}_k) = \sum_{j=1}^m A_{j,i}e^{\frac{1}{\mu}((A\mathsf{x}_k)_j - 1)} - 1 \in [-1,\infty)$.
- Perform a *mirror (descent) step* $\mathsf{z}_k \leftarrow \mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z\in\Delta_{\mathsf{box}}} \left\{\frac{1}{2}\|z - \mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k\xi_k^{(i)}, z\rangle\right\}$ for some parameter $\alpha_k \ll 1/n$ to be chosen later.
- Perform a *gradient (descent) step* $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k + \frac{1}{n\alpha_k L}(\mathsf{z}_k^{(i)} - \mathsf{z}_{k-1})$.

For this algorithm, it is not surprisingly to deduce the following theorem similar to Theorem 3.2. We include its proof to Appendix D for completeness.

Theorem 5.3. *Under the well-behavior assumption 5.1 in the covering LP problem,* $\texttt{CovLPSolver}^{\mathsf{wb}}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs some* y_T *satisfying*

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \le (1 + 4.6\varepsilon)\mathsf{OPT} \ .$$

It is now straightforward to use Markov inequality to turn the expected guarantee in Theorem 5.3 into a probabilistic one:

Corollary 5.4. *Under the well-behavior assumption 5.1 in the covering LP problem, with probability at least* $9/10$, $\texttt{CovLPSolver}^{\mathsf{wb}}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs a* $(1 + O(\varepsilon))$ *approximate solution. The expected running time is*

$$O\Big(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon} N\Big) \ .$$

Removing the well-behavior assumption. In subsequent work, after the conference presentation of this paper, Wang, Mahoney and Rao [32] showed the following theorem.

Theorem 5.5 ([32])**.** *Any covering LP with constraint matrix* $A \in \mathbb{R}^{m\times n}_{\ge 0}$ *of sparsity* N *can be converted into an equivalent but* well-behaved *covering LP with constraint matrix* $\widetilde{A} \in \mathbb{R}^{m\times n\cdot O(\log(mn/\varepsilon))}$ *and sparsity* $N \cdot O(\log(mn/\varepsilon))$*. The conversion takes time* $N \cdot O(\log(mn/\varepsilon))$*.*

As a result, we can apply our covering solver $\texttt{CovLPSolver}^{\mathsf{wb}}$ to this modified LP and apply our Theorem 5.3 to solve any covering LP in expected time $O(\frac{\log^2(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon} N)$.

6 Our Covering LP Solver in the General Case

In this section, we remove the well-behavior assumption and propose a different algorithm `CovLPSolver` that solves all covering LP instances. This new algorithm introduces a factor $1/\sqrt{\varepsilon}$ loss in the running time, but is a *direct* covering LP solver without using any reduction such as [32].

6.1 Sketching the Main Ideas

This time, we pick a random coordinate $i \in [n]$ at each iteration, and decompose $\nabla_i f(\mathsf{x}_k) = \xi + \eta$, where $\eta \in (-\infty, 0]$ is the (negative) large gradient component and $\xi \in [-\sqrt{\varepsilon}, 1]$ is the small gradient component. In other words, we intentionally truncate the gradient $\nabla_i f(\mathsf{x}_k)$ at (negative) $\sqrt{\varepsilon}$, rather than at 1 as in `PacLPSolver` (or -1 as in $\texttt{CovLPSolver}^{\mathsf{wb}}$). This is the first time our gradient truncation technique is performed non-symmetrically. Our main idea is again to perform

- a gradient (descent) step with respect to η, and
- a mirror (descent) step with respect to ξ.

The reason for this different choice of truncation threshold can be understood as follows. In `PacLPSolver`, we used Lemma 3.10 to provide a *joint* upper bound on both the ξ and the η parts of the gradient. Unfortunately, for covering LP, gradient descent only yields an upper bound on η, but not on ξ (see Lemma 6.13 later). This forces us to kill the small gradient ξ using mirror descent only, without consulting gradient descent.

If ξ were between -1 and 1 (and even if η were always zero), classical theory of mirror descent (or multiplicative weight update) only implies that the mirror step converges at a rate $\propto \varepsilon^{-2}$. For this reason, we truncate ξ a bit more and require it to reside in $[-\sqrt{\varepsilon}, 1]$. This time, using a

negative-width technique we can improve the mirror-descent convergence from ε^{-2} to $\varepsilon^{-1.5}$. Due to this weaker truncation at $-\sqrt{\varepsilon}$ instead of -1, our gradient step also converges slower, and the rate drops from $1/\varepsilon$ down to $1/\varepsilon^{1.5}$. This is precisely why we truncate the gradient at $\sqrt{\varepsilon}$, as it provides the best truncation tradeoff between gradient and mirror descent.

Another important ingredient behind our proof is a novel distance upper bound that is not commonly used in first-order analysis. Recall that, given convex function $g(x)$, traditional first-order analysis usually applies convexity and obtain upper bound $g(x) - g(x^*) \le \langle \nabla g(x), x - x^* \rangle$ on the objective distance to optimum. If $g(x) = e^{-x}$ is a univariate function, $x = -1$, and $x^* = -100$, this upper bound becomes $e^{-1} \approx e^{-1} - e^{-100} \le e^{-1} \cdot 99$, which is extremely non-tight. We replace this upper bound with a more refined one specifically designed for covering LP, which can be found in Lemma 6.8.

The detailed analysis of our covering LP solver in the general case can be found in the subsequent subsections.

6.2 Algorithm Description

We choose a slightly different starting vector from Section 5. Recall that the initial vector $x^\sharp$ has been defined in Definition 4.1. The following proposition is proved inAppendix E.

Proposition 6.1. *Letting* $x^{\mathsf{start}} = (1 + \varepsilon/2) \cdot x^\sharp + (\frac{1}{n}, \dots, \frac{1}{n})$, *we have* $\mathbb{1}^T x^{\mathsf{start}} \le 2\mathsf{OPT}'$ *and* $f_\mu(x^{\mathsf{start}}) \le 4\mathsf{OPT}$.

We study the minimization problem on $f_\mu(x)$ over the simplex

$$x \in \Delta_{\mathsf{simplex}} \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \,:\, x_i \ge 0 \ \wedge \ \mathbb{1}^T x \le 2\mathsf{OPT}'\} \ .$$

That is, over all non-negative vectors x where the coordinate sum up to at most $2\mathsf{OPT}'$. Again, this constraint $\mathbb{1}^T x \le 2\mathsf{OPT}'$ is redundant just like the Δ_{box} constraint for packing LP; however, it shall play some important role in our analysis. It is a simple fact that

Fact 6.2. $u^* = (1 + \varepsilon/2)x^* \in \Delta_{\mathsf{simplex}}$.

We also make the following choice of the thresholding function.

Definition 6.3. *The thresholding function* $\mathbb{T}^{\mathsf{c}} \colon (-\infty, 1] \to [-\beta, 1]$ *is*

$$\mathbb{T}^{\mathsf{c}}(v) \stackrel{\text{def}}{=} \begin{cases} v, & v \in [-\beta, 1]; \\ -\beta, & v < -\beta. \end{cases} \qquad \text{where } \beta \stackrel{\text{def}}{=} \sqrt{\varepsilon}$$

Like our packing solver, our algorithm `CovLPSolver` starts with the initial vector $\mathsf{x}_0 = \mathsf{y}_0 = \mathsf{z}_0 = x^{\mathsf{start}}$ introduced in Proposition 6.1 and is divided into T iterations. In each iteration, it computes a weighted midpoint $\mathsf{x}_k \leftarrow \tau \mathsf{z}_{k-1} + (1 - \tau)\mathsf{y}_{k-1}$ for some parameter $\tau \in (0, 1)$, and then proceeds to compute y_k and z_k as follows.

- Select $i \in [n]$ uniformly at random, and let $\xi_k^{(i)} = (0, \dots, 0, \mathbb{T}^{\mathsf{p}}(v), 0, \dots, 0)$ be the vector that is only non-zero at coordinate i, where $v = \nabla_i f_\mu(\mathsf{x}_k) = 1 - \sum_{j=1}^m A_{j,i} e^{\frac{1}{\mu}(1 - (A\mathsf{x}_k)_j)} \in (-\infty, 1]$.
- Perform a *mirror (descent) step* $\mathsf{z}_k \leftarrow \mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z \in \Delta_{\mathsf{simplex}}} \left\{ V_{\mathsf{z}_{k-1}}(z) + \langle (1 + \gamma) n \alpha_k \xi_k^{(i)}, z \rangle \right\}$ for some parameters $\gamma \ll 1$ and $\alpha_k \ll 1/n$, where $V_x(y) = \sum_{i=1}^n y_i \log \frac{y_i}{x_i} + x_i - y_i$ is the so-called Bregman divergence of the generalized entropy function (see Proposition 6.12 below).
- Perform a *gradient (descent) step* $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k + \delta \mathbf{e}_i$ for some value δ that is zero if $\nabla_i f_\mu(x_k) < -\beta$, and strictly positive otherwise. The precise definition of δ can be found in the pseudocode described in Algorithm 3.

Algorithm 3 $\texttt{CovLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$

Input: $A \in \mathbb{R}^{m\times n}_{\geq 0}, x^{\mathsf{start}} \in \Delta_{\mathsf{simplex}}, \varepsilon \in (0, 1/10]$.
Output: $x \in \Delta_{\mathsf{simplex}}$.

$\mu \leftarrow \frac{\varepsilon}{4\log(nm/\varepsilon)}$, $\beta \leftarrow \sqrt{\varepsilon}$, $\tau \leftarrow \frac{\mu\beta}{12n}$. ▷ parameters
$T \leftarrow \lceil \frac{1}{\tau}\log(1/\varepsilon)\rceil = O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon^{1.5}}n)$. ▷ number of iterations
$\alpha_0 \leftarrow (1-\tau)^T\frac{\varepsilon}{12n\beta}$ and $\gamma \leftarrow \frac{\varepsilon}{6\beta}$. ▷ so that $\alpha_T = \frac{\varepsilon}{12n\beta}$ and $\gamma = 2\alpha_T n$
$\mathsf{x}_0 = \mathsf{y}_0 = \mathsf{z}_0 \leftarrow x^{\mathsf{start}}$.
for $k \leftarrow 1$ **to** T **do**
    $\alpha_k \leftarrow \frac{1}{1-\tau}\alpha_{k-1}$.
    $\mathsf{x}_k \leftarrow \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$.
    Randomly select i uniformly at random from $[n]$.
    Define vector $\xi_k^{(i)}$ to be all-zero except at coordinate i, where it equals $\mathbb{T}^{\mathsf{c}}(\nabla_i f_\mu(\mathsf{x}_k))$.
        ▷ recall from (C.1) that $\nabla_i f_\mu(\mathsf{x}_k) = 1 - \sum_{j=1}^m A_{j,i}e^{\frac{1}{\mu}(1-(A\mathsf{x}_k)_j)}$
        ▷ recall from Definition 6.3 that $\mathbb{T}^{\mathsf{c}}(v) \stackrel{\text{def}}{=} \begin{cases} v, & v \in [-\beta, 1]; \\ -\beta, & v < -\beta. \end{cases}$
    $\mathsf{z}_k \leftarrow \mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z\in\Delta_{\mathsf{simplex}}}\left\{V_{\mathsf{z}_{k-1}}(z) + \langle(1+\gamma)n\alpha_k\xi_k^{(i)}, z\rangle\right\}$. ▷ See Proposition 6.12
    if $\nabla_i f_\mu(\mathsf{x}_k) < -\beta$ **then**
        Denote by π the permutation that satisfies $A_{\pi(1),i} \leq \cdots \leq A_{\pi(m),i}$.
        Pick $j^* \in [m]$ such that $\begin{cases} \sum_{j<j^*} A_{\pi(j),i}\cdot p_{\pi(j)}(\mathsf{x}_k) < 1+\beta \\ \sum_{j\leq j^*} A_{\pi(j),i}\cdot p_{\pi(j)}(\mathsf{x}_k) \geq 1+\beta \end{cases}$.
            ▷ Such a $j^* \in [m]$ must exist because $\sum_{j=1}^m A_{ji}\cdot p_j(\mathsf{x}_k) \geq 1+\beta$.
        $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k + \delta\cdot\mathbf{e}_i$ where $\delta = \frac{\mu\beta}{2A_{\pi(j^*),i}}$.
    else
        $\mathsf{y}_k \leftarrow \mathsf{y}_k^{(i)} \stackrel{\text{def}}{=} \mathsf{x}_k$.
    end if
end for
return y_T.

We use the superscript $^{(i)}$ on $\xi_k^{(i)}$, $\mathsf{y}_k^{(i)}$ and $\mathsf{z}_k^{(i)}$ to emphasize that the value depends on the choice of i. We have used generic parameters τ, α_k, T in the above description and their precise values are presented in Algorithm 3.

6.3 Convergence Statement

The main convergence theorem of this section is as follows:

Theorem 6.4. $\texttt{CovLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs some* y_T *satisfying*

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \leq (1+9\varepsilon)\mathsf{OPT} \ .$$

The bound on the running time is due to the following lemma, which is analogous to its packing counterpart and is proved in Section G.

Lemma 6.5. *Each iteration of* $\texttt{CovLPSolver}$ *can be implemented to run in expected* $O(N/n)$ *time. The total expected running time is* $O(TN/n)$.

It is again straightforward to use Markov inequality to turn the expected guarantee in Theorem 6.4 into a probabilistic one:

Corollary 6.6. *With probability at least* 9/10, $\texttt{CovLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs a* $(1+O(\varepsilon))$ *approximate solution to the covering LP program. The expected running time is* $O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon^{1.5}}N)$.

Before delving into the proof of Theorem 6.4, we make the following observations. Since the x^{start} satisfies $\mathbb{1}^T x^{\mathsf{start}} \le 2\mathsf{OPT}'$ by Proposition 6.1, we have $\mathsf{z}_0 = x^{\mathsf{start}} \in \Delta_{\mathsf{simplex}}$. Also, the mirror descent step ensures $\mathsf{z}_{k,i} > 0$ for all rounds k and coordinates i, as well as $\mathsf{z}_k \in \Delta_{\mathsf{simplex}}$ for all rounds k. However, we note that x_k and y_k may not necessarily lie inside $\Delta_{\mathsf{simplex}}$, but will always stay non-negative. We summarize these properties as follows:

$$\forall k \in \{0, 1, \ldots, T\}, \qquad \mathsf{x}_k, \mathsf{y}_k \ge 0, \quad \mathsf{z}_k > 0, \quad \mathsf{z}_k \in \Delta_{\mathsf{simplex}} \ .$$

We are now ready to prove Theorem 6.4 in the subsequent subsections.

6.4 Step 1: Distance Adjustment

Classically, using the convexity argument one can obtain $f_\mu(\mathsf{x}_k) - f_\mu(u) \le \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u\rangle$ for every $u \in \Delta_{\mathsf{simplex}}$. In particular, if u is the optimal point, the right hand side is a simple upper bound on the objective distance from the current point $f_\mu(\mathsf{x}_k)$ to the optimum. This simple upper bound is used by all classical first-order analyses.

In this subsection, we strengthen this upper bound in the special case of $u = u^* \stackrel{\text{def}}{=} (1+\varepsilon/2)x^*$. Define $\widetilde{A}$ be the *adjusted matrix* of A described as follows.

Definition 6.7 (Adjusted matrix $\widetilde{A}$)**.** *For each row* $j \in [m]$*, if* $(Au^*)_j \le 2$ *then we keep this row and define* $\widetilde{A}_{j:} \stackrel{\text{def}}{=} A_{j:}$*. Otherwise, —that is, if* $(Au^*)_j > 2$*— we define* $\widetilde{A}_{j:} \stackrel{\text{def}}{=} \frac{2}{(Au^*)_j} \cdot A_{j:}$ *to be the same* j*-th row* $A_{j:}$*, but scaled down by a factor of* $\frac{2}{(Au^*)_j}$*. It is clear from this definition that*

$$A_{ji} \ge \widetilde{A}_{ji} \text{ for all } (i,j) \in [n] \times [m] \quad \text{and} \quad (1+\varepsilon)\mathbb{1} \le \widetilde{A}u^* \le 2\mathbb{1}.$$

We now state the following stronger upper bound:

Lemma 6.8 (Distance Adjustment)**.**

$$\begin{aligned} f_\mu(\mathsf{x}_k) - f_\mu(u^*) &\le \langle \mathbb{1} - A^T p(\mathsf{x}_k), \mathsf{x}_k - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle + \varepsilon\mathsf{OPT} \\ &= \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle + \varepsilon\mathsf{OPT} \end{aligned}$$

At high level, ignoring the negligible term $\varepsilon\mathsf{OPT}$ on the right hand side, the above upper bound strengthens the classical bound due to the extra term of $\langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle$. This extra term is always non-positive since $\widetilde{A} \le A$ coordinate-wisely, but may be very negative in certain cases.

The intuition behind the proof is to realize that the convexity inequality $e^b - e^a \le \langle e^b, b-a\rangle$ on the exponential function becomes far from tight when $a \ll 0$. For instance, when $b = 2$ and $a = -10$, we have $e^2 - e^{-10} \le 12e^2$; when $b = 2$ and $a = -100$, we only get $e^2 - e^{-100} \le 102e^2$. Although $e^{-100} \approx e^{-10}$, the two upper bounds are off from each other by a factor of 10. Therefore, when necessary, we can 'elevate' a to some higher value in order to obtain a tighter upper bound. We defer the detailed proof to Appendix E.

6.5 Step 2: Gradient Truncation

Let us separate the indices $i \in [n]$ into large and small ones.

Definition 6.9. *We make the following definitions.*

- *Let* $\xi_k \in [-\beta, 1]^n$ *be the* truncated gradient *so that* $\xi_{k,i} = \mathbb{T}^{\mathsf{c}}(\nabla_i f_\mu(x_k))$ *for each* $i \in [n]$.

- *Let* $B_k \stackrel{\text{def}}{=} \{i \in [n] : \xi_{k,i} \neq \nabla_i f_\mu(\mathsf{x}_k)\}$ *be the set of* large indices.
- *Let* $\eta_k \in (-\infty, 0]^n$ *be the* large gradient *so that* $\nabla f_\mu(\mathsf{x}_k) = \xi_k + \eta_k$. *It is clear that*
$$\eta_{k,i} = 0 \text{ for every } i \notin B, \text{ and } \eta_{k,i} = (1+\beta) - (A^T p(\mathsf{x}_k))_i \text{ for every } i \in B.$$
- *Let* $\widetilde{\eta}_k \in (-\infty, \infty)^n$ *be the* adjusted large gradient *so that*
$$\widetilde{\eta}_{k,i} = 0 \text{ for every } i \notin B, \text{ and } \widetilde{\eta}_{k,i} = (1+\beta) - (\widetilde{A}^T p(\mathsf{x}_k))_i \text{ for every } i \in B.$$

For the rest of this section, we denote by $\eta_k^{(i)} = (0, \ldots, 0, \eta_{k,i}, 0, \ldots, 0)$, *the vector that is zero at all coordinates other than* i, *and equals to* $\eta_{k,i}$ *at location* i. *We similarly define* $\xi_k^{(i)}$ *as well as* $\widetilde{\eta}_k^{(i)}$.

We state the following key lemma that is very analogous to (3.1) in the packing LP case. Note that if one uses $\eta_k^{(i)}$ instead of $\widetilde{\eta}_k^{(i)}$, the proof becomes identical to that of (3.1). The reason that we can replace $\eta_k^{(i)}$ with $\widetilde{\eta}_k^{(i)}$ —thus giving a better upper bound— is precisely due to the distance adjustment introduced in Lemma 6.8.

Lemma 6.10 (distance upper bound)**.**

$$f_\mu(\mathsf{x}_k) - f_\mu(u^*) \leq \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \mathbf{E}_i\Big[\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\rangle\Big] + \mathbf{E}_i\Big[\langle n\widetilde{\eta}_k^{(i)}, -u^*\rangle\Big] + \varepsilon\mathsf{OPT} \ .$$

The proof of the above lemma is a simple repetition of that of (3.1), but replacing the classical distance upper bound with our adjusted one. See Appendix E for details.

6.6 Step 3: Mirror Descent Guarantee

Our update $\mathsf{z}_k^{(i)} \stackrel{\text{def}}{=} \arg\min_{z \in \Delta_{\mathsf{simplex}}} \left\{V_{\mathsf{z}_{k-1}}(z) + \langle (1+\gamma) n\alpha_k \xi_k^{(i)}, z\rangle\right\}$ is, by its definition, a mirror descent step. We begin by explaining an attempt that is too weak for obtaining the $\varepsilon^{-1.5}$ convergence rate.

Using the classical theory of mirror descent, it is not hard to repeat the proof of Lemma 3.6 —although changing the distance function from $\|\cdot\|_A^2$ to $V_x(y)$— and obtain that, for every $u \in \Delta_{\mathsf{simplex}}$,

$$\mathbf{E}_i\big[\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u\big\rangle\big] \leq V_{\mathsf{z}_{k-1}}(u) - \mathbf{E}_i\Big[V_{\mathsf{z}_k^{(i)}}(u)\Big] + O(\alpha_k^2 n)\mathsf{OPT} \ .$$

The above inequality can be made true whenever ξ_i is between -1 and 1 for each coordinate i, but only yields the known ε^{-2} convergence rate. Here, ± 1 is also know as the *width* from multiplicative-weight-update languages [6].

Fortunately, since we have required ξ_i to be only between $-\beta$ and 1, the $O(\alpha_k^2 n)$ factor can essentially be improved to $O(\alpha_k^2 \beta n)$. This is an improvement whenever $\beta \ll 1$, and we call it the *negative-width technique.*[13]

Formally, we prove that

Lemma 6.11 (mirror descent)**.** *Denoting by* $\gamma \stackrel{\text{def}}{=} 2\alpha_T n$, *we have*

$$\mathbf{E}_i\big[\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\big\rangle\big] \leq V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - \mathbf{E}_i\Big[V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big)\Big] + 12\mathsf{OPT} \cdot \gamma\alpha_k\beta \ .$$

[13] This negative width technique is related to [6, Definition 3.2], where the authors analyze the classical multiplicative weight update method in a special case when the oracle returns loss values only between $-\ell$ and ρ, for $\ell \ll \rho$. This technique is in fact related to a more general theory of mirror descent, known as the local-norm convergence, that we have summarized in a separate paper [1].

The proof can be found in Appendix E.

Although defined in a variational way, it is perhaps beneficial to explicitly describe how to implement this mirror step. The following proposition is straightforward but anyways proved in Appendix E:

Proposition 6.12. *If $\mathsf{z}_{k-1} \in \Delta_{\mathsf{simplex}}$ and $\mathsf{z}_{k-1} > 0$, the minimizer $z = \arg\min_{z \in \Delta_{\mathsf{simplex}}} \{V_{\mathsf{z}_{k-1}}(z) + \langle \delta \mathbf{e}_i, z\rangle\}$ for any scalar $\delta \in \mathbb{R}$ and basis vector $\mathbf{e}_i$ can be computed as follows:*

1. *$z \leftarrow \mathsf{z}_{k-1}$.*
2. *$z_i \leftarrow z_i \cdot e^{-\delta}$.*
3. *If $\mathbb{1}^T z > 2\mathsf{OPT}'$, $z \leftarrow \frac{2\mathsf{OPT}'}{\mathbb{1}^T z} z$.*
4. *Return z.*

6.7 Step 4: Gradient Descent Guarantee

We claim that our gradient step $\mathsf{x}_k \to \mathsf{y}_k^{(i)}$ never increases the objective for all choices of i. In addition, it decreases the objective by an amount proportional to the adjusted large gradient $\widetilde{\eta}_k^{(i)}$.

Lemma 6.13 (gradient descent)**.** *For every $i \in [n]$, we have*

(a) $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq 0$, and

(b) $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq \frac{\mu\beta}{12} \cdot \langle -\widetilde{\eta}_k^{(i)}, u^* \rangle$.

The proof of Lemma 6.13 is quite technical and can be found in Appendix E.

At high level, one would generally hope to prove that the gradient step decreases the objective by an amount proportional to the large gradient $\eta_k^{(i)}$, rather than the adjusted large gradient $\widetilde{\eta}_k^{(i)}$. If that were true, the entire proof structure of our covering LP convergence would become much closer to that of packing LP, and there would be absolutely no need for the introduction of the distance adjustment in Section 6.4, as well as the definitions of $\widetilde{A}$ and $\widetilde{\eta}$.

Unfortunately, if one replaces $\widetilde{\eta}$ with η in the above lemma, the inequality is *far* from being correct. The reason behind it is very similar to that we have summarized in Section 6.4, and related to the unpleasant behavior of the exponential penalty function.

6.8 Step 5: Putting All Together

Combining Lemma 6.10, Lemma 6.11, and Lemma 6.13, we obtain that

$$\begin{aligned}
&\alpha_k\big(f_\mu(\mathsf{x}_k) - f_\mu(u^*)\big) - \alpha_k \varepsilon \mathsf{OPT} \\
\leq\ & \frac{(1-\tau)\alpha_k}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \mathbf{E}_i\Big[\alpha_k \langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\rangle\Big] + \mathbf{E}_i\Big[\alpha_k \langle n\widetilde{\eta}_k^{(i)}, -u^*\rangle\Big] \\
\leq\ & \frac{(1-\tau)\alpha_k}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - \mathbf{E}_i\Big[V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big)\Big] \\
&\qquad + 12\mathsf{OPT} \cdot \gamma\alpha_k\beta + \mathbf{E}_i\Big[\frac{12\alpha_k n}{\mu\beta}\big(f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})\big)\Big]
\end{aligned}$$

Remark 6.14. Above, the quantity "$12\mathsf{OPT} \cdot \gamma\alpha_k\beta$" is the loss term introduced by the mirror descent. Unlike the packing LP case —see (3.2)— this loss term is not dominated by the gradient step. (If one could do so, this would turn our `CovLPSolver` into an ε^{-1} convergence rate.)

The quantity "$\alpha_k\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\rangle$" is the loss introduced by the (adjusted) large gradient $\widetilde{\eta}$, and is dominated by our gradient step progress owing to Lemma 6.13. This is similar to the packing LP case —see Lemma 3.10.

From here, let us use the special choice of $\tau = \frac{\mu\beta}{12n}$. We obtain that

$$\begin{aligned}
&-\alpha_k\big(f_\mu(u^*) + \varepsilon\mathsf{OPT}\big) \\
&\le 12\gamma\alpha_k\beta\mathsf{OPT} + \frac{(1-\tau)\alpha_k}{\tau} f_\mu(\mathsf{y}_{k-1}) + V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - \mathbf{E}_i\Big[\frac{\alpha_k}{\tau} f_\mu(\mathsf{y}_k^{(i)}) + V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big)\Big] \ .
\end{aligned}$$

Use the choice $\alpha_k = \frac{\alpha_{k-1}}{1-\tau}$ and telescoping the above inequality for $k = 1, \dots, T$, we have

$$-\big(\sum_{k=1}^{T}\alpha_k\big)\big(f_\mu(u^*) + \varepsilon\mathsf{OPT}\big) \le \big(\sum_{k=1}^{T}\alpha_k\big)\cdot 12\gamma\beta\mathsf{OPT} + \frac{\alpha_0}{\tau} f_\mu(\mathsf{y}_0) + V_{\mathsf{z}_0}\big(\frac{u^*}{1+\gamma}\big) - \frac{\alpha_T}{\tau}\mathbf{E}\big[f_\mu(\mathsf{y}_T)\big] \ .$$

We compute that $\sum_{k=1}^{T}\alpha_k = \alpha_T \cdot \sum_{k=0}^{T-1}(1-\tau)^k = \alpha_T \cdot \frac{1-(1-\tau)^T}{\tau} < \frac{\alpha_T}{\tau}$, and recall that $\gamma = 2\alpha_T n$. Therefore, we rearrange and get

$$\begin{aligned}
&\frac{\alpha_T}{\tau}\mathbf{E}\big[f_\mu(\mathsf{y}_T)\big] \le \frac{\alpha_T}{\tau}\big(f_\mu(u^*) + \varepsilon\mathsf{OPT}\big) + \frac{\alpha_T}{\tau}\cdot 12\gamma\beta\mathsf{OPT} + \frac{\alpha_0}{\tau} f_\mu(\mathsf{y}_0) + V_{\mathsf{z}_0}\big(\frac{u^*}{1+\gamma}\big) \ , \\
&\implies \mathbf{E}\big[f_\mu(\mathsf{y}_T)\big] \le f_\mu(u^*) + \varepsilon\mathsf{OPT} + 24\alpha_T n\beta\mathsf{OPT} + (1-\tau)^T f_\mu(\mathsf{y}_0) + \frac{\tau}{\alpha_T} V_{\mathsf{z}_0}\big(\frac{u^*}{1+\gamma}\big) \ . \qquad (6.1)
\end{aligned}$$

From this point, we need to use our special choice of the initial point $\mathsf{x}_0 = \mathsf{y}_0 = \mathsf{z}_0 = x^{\mathsf{start}}$ (see Proposition 6.1), which implies that $f_\mu(\mathsf{y}_0) \le 4\mathsf{OPT}$ and $\mathbb{1}^T x^{\mathsf{start}} \le 4\mathsf{OPT}$. We also have

$$\begin{aligned}
V_{\mathsf{z}_0}\big(\frac{u^*}{1+\gamma}\big) = V_{x^{\mathsf{start}}}\big(\frac{u^*}{1+\gamma}\big) &= \sum_{i=1}^{n} \frac{u_i^*}{1+\gamma}\log\frac{u_i^*}{(1+\gamma)x_i^{\mathsf{start}}} + x_i^{\mathsf{start}} - \frac{u_i^*}{1+\gamma} \\
&\overset{①}{\le} \sum_{i=1}^{n} u_i^*\log(u_i^*\cdot n) + 4\mathsf{OPT} \overset{②}{\le} (2\log(nm)+4)\cdot\mathsf{OPT} \ .
\end{aligned}$$

Above, inequality ① follows because $x_i^{\mathsf{start}} \ge 1/n$ for all $i \in [n]$ according to the definition in Proposition 6.1; inequality ② follows because $u_i^* \le (1+\varepsilon/2)x_i^* \le (1+\varepsilon/2)\mathsf{OPT} \le (1+\varepsilon/2)m$ and $\mathbb{1}^T u_i^* = (1+\varepsilon/2)\mathsf{OPT}$, as well as the fact that ε is sufficiently small.

Finally, we choose $\beta = \sqrt{\varepsilon}$, $\alpha_T = \frac{\varepsilon}{12n\beta}$, and $T = \lceil\frac{1}{\tau}\log(1/\varepsilon)\rceil$. Substituting into (6.1) all of these parameters, along with the aforementioned inequalities $f_\mu(\mathsf{y}_0) \le 4\mathsf{OPT}$ and $V_{\mathsf{z}_0}\big(\frac{u^*}{1+\gamma}\big) \le (2\log(nm)+4)\cdot\mathsf{OPT}$, as well as $f_\mu(u^*) \le (1+\varepsilon)\mathsf{OPT}$ from Proposition 4.3.b, we obtain that

$$\mathbf{E}\big[f_\mu(\mathsf{y}_T)\big] \le (1+\varepsilon)\mathsf{OPT} + \varepsilon\mathsf{OPT} + 2\varepsilon\mathsf{OPT} + \varepsilon f_\mu(\mathsf{y}_0) + \frac{\mu\beta/12n}{\varepsilon/12n\beta}(2\log(nm)+4)\mathsf{OPT} = (1+9\varepsilon)\mathsf{OPT} \ .$$

This finishes the proof of Theorem 6.4. □

Appendix

A Missing Proofs for Section 2

Proposition 2.3. *Let $\mu = \frac{\varepsilon}{4\log(nm/\varepsilon)}$ and x^* be an optimal solution for the packing LP (1.1). Then:*

(a) $f_\mu(u^*) \le -(1-\varepsilon)\mathsf{OPT}$ for $u^* \stackrel{\text{def}}{=} (1-\varepsilon/2)x^* \in \Delta_{\mathsf{box}}$.
(b) $f_\mu(x) \ge -(1+\varepsilon)\mathsf{OPT}$ for every $x \in \Delta_{\mathsf{box}}$.
(c) If $x \in \Delta_{\mathsf{box}}$ satisfies $f_\mu(x) \le -(1-O(\varepsilon))\mathsf{OPT}$, then $\frac{1}{1+\varepsilon}x$ is a $(1-O(\varepsilon))$-approximate solution to the packing LP.

Proof.

(a) We have $\mathbb{1}^T u^* = (1-\varepsilon/2)\mathsf{OPT}$ by the definition of OPT. Also, from the feasibility constraint $Ax^* \le \mathbb{1}$ in the packing LP, we have $Au^* - \mathbb{1} \le -\varepsilon/2 \cdot \mathbb{1}$, and can compute $f_\mu(u^*)$ as follows:

$$\begin{aligned} f_\mu(u^*) = \mu \sum_j e^{\frac{1}{\mu}((Au^*)_j - 1)} - \mathbb{1}^T u^* &\le \mu \sum_j e^{\frac{-\varepsilon/2}{\mu}} - (1-\varepsilon/2)\mathsf{OPT} \\ &\le \frac{\mu m}{(nm)^2} - (1-\varepsilon/2)\mathsf{OPT} \le -(1-\varepsilon)\mathsf{OPT} \ . \end{aligned}$$

(b) Suppose towards contradiction that $f_\mu(x) < -(1+\varepsilon)\mathsf{OPT}$. Since $f_\mu(x) > -\mathbb{1}^T x$, it must satisfy that $\mathbb{1}^T x > (1+\varepsilon)\mathsf{OPT}$. Suppose that $\mathbb{1}^T x = (1+v)\mathsf{OPT}$ for some $v > \varepsilon$. By the definition of OPT, we must have that $Ax < (1+v)\mathbb{1}$ is broken, and therefore there exists some $j \in [m]$ satisfying that $(Ax)_j \ge 1+v$. In such a case, the objective

$$f_\mu(x) \ge \mu e^{v/\mu} - (1+v)\mathsf{OPT} = \frac{\varepsilon}{4\log(nm)}\Big((\frac{nm}{\varepsilon})^4\Big)^{v/\varepsilon} - (1+v)\mathsf{OPT} \ge \Big(\Big((\frac{nm}{\varepsilon}\Big)^2)^{v/\varepsilon} - (1+v)\Big)\mathsf{OPT} > 0$$

giving a contradiction to the assumption that $f_\mu(x) < 0$.

(c) Suppose x satisfies $f_\mu(x) \le -(1-O(\varepsilon))\mathsf{OPT} \le 0$ and we first want to show $Ax \le (1+\varepsilon)\mathbb{1}$. Let us assume that $v = \max_j((Ax)_j - 1) \ge 0$ because otherwise we will have $Ax \le \mathbb{1}$. Under this definition, we have $Ax \le (1+v)\mathbb{1}$ and therefore $\mathbb{1}^T x \le (1+v)\mathsf{OPT}$ by the definition of OPT. We compute $f_\mu(x)$ as follows.

$$f_\mu(x) \ge \mu e^{\frac{v}{\mu}} - (1+v)\mathsf{OPT} \ge \mu\Big((\frac{nm}{\varepsilon})^4\Big)^{v/\varepsilon} - (1+v)n = \frac{\varepsilon}{4\log(nm)}\Big((\frac{nm}{\varepsilon})^4\Big)^{v/\varepsilon} - (1+v)n \ .$$

It is easy to see that the above quantity is positive whenever $v \ge \varepsilon$, and therefore, to satisfy $f_\mu(x) \le 0$ we must have $v \le \varepsilon$, which is equivalent to $Ax \le (1+\varepsilon)\mathbb{1}$.

Next, because $-\mathbb{1}^T x \le f_\mu(x) \le -(1-O(\varepsilon))\mathsf{OPT}$, we know that x yields an objective $\mathbb{1}^T x \ge (1-O(\varepsilon))\mathsf{OPT}$. Letting $x' = \frac{1}{1+\varepsilon}x$, we both have that x' is feasible (i.e., $Ax' \le \mathbb{1}$), and x' has an objective $\mathbb{1}^T x'$ at least as large as $(1-O(\varepsilon))\mathsf{OPT}$. □

Lemma 2.6. *Define the smoothness parameter $L \stackrel{\text{def}}{=} \frac{4}{\mu}$. Then, for every $x \ge 0$, every $i \in [n]$, and every $\lambda \le \frac{1}{L\|A_{:i}\|_\infty}$:*

(a) If $|\nabla_i f_\mu(x)| \le 1$, then $\big|\nabla_i f_\mu(x + \lambda \mathbf{e}_i) - \nabla_i f_\mu(x)\big| \le L\|A_{:i}\|_\infty \cdot |\lambda|$.

(b) If $\nabla_i f_\mu(x) \geq 1$, *then* $\nabla_i f_\mu(x+\lambda \mathbf{e}_i) \geq \Big(1 - \frac{\|A_{:i}\|_\infty L}{2}|\lambda|\Big)\nabla_i f_\mu(x)$.

Proof of Lemma 2.6. Using the fact that $\nabla_i f_\mu(x) > -1$ for all x, we have:

$$\begin{aligned}\Big|\log \frac{\nabla_i f_\mu(x+\lambda\mathbf{e}_i)+1}{\nabla_i f_\mu(x)+1}\Big| &\overset{①}{=} \Big|\int_0^\lambda \frac{\nabla^2_{ii} f_\mu(x+\nu\mathbf{e}_i)}{\nabla_i f_\mu(x+\nu\mathbf{e}_i)+1} d\nu\Big| \\ &\overset{②}{=} \frac{1}{\mu}\Big|\int_0^\lambda \frac{(A^T \mathrm{diag}\{p(x+\nu\mathbf{e}_i)\}A)_{ii}}{(A^T p(x+\nu\mathbf{e}_i))_i} d\nu\Big| \\ &\overset{③}{\leq} \frac{\|A_{:i}\|_\infty}{\mu}|\lambda| \overset{④}{=} \frac{\|A_{:i}\|_\infty L}{4}|\lambda| \ .\end{aligned}$$

Above, ① holds because $\int_0^\lambda g'(\nu)d\nu = g(\lambda) - g(0)$; ② holds according to Fact 2.5; ③ is because the numerator is $\sum_j A_{j,i}^2 p_j$ while the denominator is $\sum_j A_{j,i} p_j$; ④ holds because $L = \frac{4}{\mu}$. This immediately implies the following multiplicative bound:

$$e^{-\frac{\|A_{:i}\|_\infty L}{4}|\lambda|} \leq \frac{\nabla_i f_\mu(x+\lambda\mathbf{e}_i)+1}{\nabla_i f_\mu(x)+1} \leq e^{\frac{\|A_{:i}\|_\infty L}{4}|\lambda|}.$$

Our assumption on λ implies $\frac{\|A_{:i}\|_\infty L}{4}|\lambda| \leq \frac{1}{4}$, so that we can use the approximation $x \leq e^x - 1 \leq 1.2x$ over $x \in [-\frac{1}{4}, \frac{1}{4}]$. This yields the simpler bound:

$$-\frac{\|A_{:i}\|_\infty L}{4}|\lambda| \leq \frac{\nabla_i f_\mu(x+\lambda\mathbf{e}_i) - \nabla_i f_\mu(x)}{\nabla_i f_\mu(x)+1} \leq 1.2\frac{\|A_{:i}\|_\infty L}{4}|\lambda|.$$

(a) Assuming that $\nabla_i f_\mu(x) \in (-1, 1]$, we have:

$$\Big|\nabla_i f_\mu(x+\lambda\mathbf{e}_i) - \nabla_i f_\mu(x)\Big| \leq 2.4 \cdot \frac{\|A_{:i}\|_\infty L}{4}|\lambda| \leq \|A_{:i}\|_\infty L|\lambda| \ .$$

(b) Assuming $\nabla_i f_\mu(x) \geq 1$, we have

$$\nabla_i f_\mu(x+\lambda\mathbf{e}_i) \geq \nabla_i f_\mu(x) - \frac{\|A_{:i}\|_\infty L}{4}|\lambda|\Big(\nabla_i f_\mu(x)+1\Big) \geq \Big(1 - \frac{\|A_{:i}\|_\infty L}{2}|\lambda|\Big)\nabla_i f_\mu(x) \ . \quad \square$$

Fact 2.7. *Defining* $x_i^{\mathsf{start}} \overset{\text{def}}{=} \frac{1-\varepsilon/2}{n\|A_{:i}\|_\infty}$ *for for each* $i \in [n]$, *we have* $x^{\mathsf{start}} \in \Delta_{\mathsf{box}}$ *and* $f_\mu(x^{\mathsf{start}}) \leq -\frac{1-\varepsilon}{n}$.

Proof. Using the fact that $Ax^{\mathsf{start}} - \mathbb{1} \leq -\varepsilon/2 \cdot \mathbb{1}$, we compute $f_\mu(x^{\mathsf{start}})$ as follows:

$$\begin{aligned}f_\mu(x^{\mathsf{start}}) &= \mu\sum_j e^{\frac{1}{\mu}((Ax^{\mathsf{start}})_j - 1)} - \mathbb{1}^T x^{\mathsf{start}} \leq \mu\sum_j e^{\frac{-\varepsilon/2}{\mu}} - \frac{1-\varepsilon/2}{n} \\ &\leq \frac{\mu m}{(nm)^2} - \frac{1-\varepsilon/2}{n} \leq -\frac{1-\varepsilon}{n} \ .\end{aligned}$$

Above, we have used that $\mathbb{1}^T x^{\mathsf{start}} \geq x_i^{\mathsf{start}} = \frac{1-\varepsilon/2}{n}$, where i is the column such that $\|A_{:i}\|_\infty = 1$. $\square$

B Missing Proofs for Section 3

Lemma 3.5. *We have* $\mathsf{x}_k, \mathsf{y}_k, \mathsf{z}_k \in \Delta_{\mathsf{box}}$ *for all* $k = 0, 1, \ldots, T$.

Proof. This is true at the beginning as $\mathsf{x}_0 = \mathsf{y}_0 = x^{\mathsf{start}} \in \Delta_{\mathsf{box}}$ (see Fact 2.7) and $\mathsf{z}_0 = 0 \in \Delta_{\mathsf{box}}$.

In fact, it suffices for us to show that for every $k \geq 0$, $\mathsf{y}_k = \sum_{l=0}^{k} \gamma_k^l \mathsf{z}_l$ for some scalers γ_k^l satisfying $\sum_l \gamma_k^l = 1$ and $\gamma_k^l \geq 0$ for each $l = 0, \ldots, k$. If this is true, we can prove the lemma by induction: at each iteration k,

1. $\mathsf{x}_k = \tau \mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$ must be in Δ_{box} because y_{k-1} and z_{k-1} are and $\tau \in [0,1]$,
2. z_k is in Δ_{box} by the definition that $\mathsf{z}_k = \arg\min_{z \in \Delta_{\mathsf{box}}}\{\cdots\}$, and
3. y_k is also in Δ_{box} because $\mathsf{y}_k = \sum_{l=0}^{k} \gamma_k^l \mathsf{z}_l$ is a convex combination of the z_l's and Δ_{box} is convex.

For the rest of the proof, we only need to show that $\mathsf{y}_k = \sum_{l=0}^{k} \gamma_k^l \mathsf{z}_l$ for[14]

$$\gamma_k^l = \begin{cases} (1-\tau)\gamma_{k-1}^l, & l = 0, \ldots, k-2; \\ \left(\frac{1}{n\alpha_{k-1}L} - \frac{1}{n\alpha_k L}\right) + \tau\left(1 - \frac{1}{n\alpha_{k-1}L}\right), & l = k-1; \\ \frac{1}{n\alpha_k L}, & l = k. \end{cases}$$

This is true at the base case because $\alpha_0 = \frac{1}{nL}$. It is also true at $k = 1$ because $\mathsf{y}_1 = \mathsf{x}_1 + \frac{1}{n\alpha_1 L}(\mathsf{z}_1 - \mathsf{z}_0) = \frac{1}{n\alpha_1 L}\mathsf{z}_1 + \left(1 - \frac{1}{n\alpha_1 L}\right)\mathsf{z}_0$. For the general k, we have

$$\begin{aligned} \mathsf{y}_k &= \mathsf{x}_k + \frac{1}{n\alpha_k L}(\mathsf{z}_k - \mathsf{z}_{k-1}) \\ &= \tau \mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1} + \frac{1}{n\alpha_k L}(\mathsf{z}_k - \mathsf{z}_{k-1}) \\ &= \tau \mathsf{z}_{k-1} + (1-\tau)\Big(\sum_{l=0}^{k-2} \gamma_{k-1}^l \mathsf{z}_l + \frac{1}{n\alpha_{k-1}L}\mathsf{z}_{k-1}\Big) + \frac{1}{n\alpha_k L}(\mathsf{z}_k - \mathsf{z}_{k-1}) \\ &= \Big(\sum_{l=0}^{k-2}(1-\tau)\gamma_{k-1}^l \mathsf{z}_l\Big) + \left(\Big(\frac{1}{n\alpha_{k-1}L} - \frac{1}{n\alpha_k L}\Big) + \tau\Big(1 - \frac{1}{n\alpha_{k-1}L}\Big)\right)\mathsf{z}_{k-1} + \frac{1}{n\alpha_k L}\mathsf{z}_k \ . \end{aligned}$$

Therefore, we obtain $\mathsf{y}_k = \sum_{l=0}^{k} \gamma_k^l \mathsf{z}_l$ as desired.

It is now easy to check that under our definition of α_k (which satisfies $\alpha_k \geq \alpha_{k-1}$ and $\alpha_k \geq \alpha_0 = \frac{1}{nL}$, we must have $\gamma_k^l \geq 0$ for all k and l. Also,

$$\begin{aligned} \sum_l \gamma_k^l &= \sum_{l=0}^{k-2}(1-\tau)\gamma_{k-1}^l + \left(\Big(\frac{1}{n\alpha_{k-1}L} - \frac{1}{n\alpha_k L}\Big) + \tau\Big(1 - \frac{1}{n\alpha_{k-1}L}\Big)\right) + \frac{1}{n\alpha_k L} \\ &= (1-\tau)\Big(1 - \frac{1}{n\alpha_{k-1}L}\Big) + \left(\Big(\frac{1}{n\alpha_{k-1}L} - \frac{1}{n\alpha_k L}\Big) + \tau\Big(1 - \frac{1}{n\alpha_{k-1}L}\Big)\right) + \frac{1}{n\alpha_k L} = 1 \ . \end{aligned}$$

□

Lemma 3.6. *When* $\mathsf{z}_k^{(i)} = \arg\min_{z \in \Delta_{\mathsf{box}}}\left\{\frac{1}{2}\|z - \mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k \xi_k^{(i)}, z\rangle\right\}$*, we have for every* $u \in \Delta_{\mathsf{box}}$

$$\left\langle n\alpha_k \xi_k^{(i)}, \mathsf{z}_{k-1} - u\right\rangle \leq n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)}\right\rangle + \frac{1}{2}\|\mathsf{z}_{k-1} - u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)} - u\|_A^2 \ .$$

Proof. Denoting by $V_a(b) = \frac{1}{2}\|b - a\|_A^2$ as a function of $b \in \Delta_{\mathsf{box}}$ parameterized at $a \in \Delta_{\mathsf{box}}$, we have that $\nabla_i V_a(b) = \|A_{:i}\|_\infty \cdot (a_i - b_i)$. In the optimization language, $V_a(b)$ is also known as the Bregman divergence of the $\|\cdot\|_A^2$ regularizer.

[14] We wish to point out that this proof coincides with a lemma from the accelerated coordinate descent theory of Fercoq and Richtárik [16]. Their paper is about optimizing an objective function that is Lipschitz smooth, and thus irrelevant to our work.

We deduce the following sequence of inequalities:

$$\begin{aligned}
\big\langle n\alpha_k\xi_k^{(i)}, \mathsf{z}_{k-1}-u\big\rangle &= \big\langle n\alpha_k\xi_k^{(i)}, \mathsf{z}_{k-1}-\mathsf{z}_k^{(i)}\big\rangle + \big\langle n\alpha_k\xi_k^{(i)}, \mathsf{z}_k^{(i)}-u\big\rangle \\
&\overset{①}{\le} \big\langle n\alpha_k\xi_k^{(i)}, \mathsf{z}_{k-1}-\mathsf{z}_k^{(i)}\big\rangle + \big\langle -\nabla V_{\mathsf{z}_{k-1}}(\mathsf{z}_k^{(i)}), \mathsf{z}_k^{(i)}-u\big\rangle \\
&\overset{②}{=} \big\langle n\alpha_k\xi_k^{(i)}, \mathsf{z}_{k-1}-\mathsf{z}_k^{(i)}\big\rangle - \frac{1}{2}\|\mathsf{z}_{k-1}-\mathsf{z}_k^{(i)}\|_A^2 + \frac{1}{2}\|\mathsf{z}_{k-1}-u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)}-u\|_A^2 \\
&\overset{③}{=} n^2\alpha_k^2 L\Big(\big\langle \xi_k^{(i)}, \mathsf{x}_k-\mathsf{y}_k\big\rangle - \frac{L}{2}\|\mathsf{x}_k-\mathsf{y}_k\|_A^2\Big) + \frac{1}{2}\|\mathsf{z}_{k-1}-u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)}-u\|_A^2 \\
&\le n^2\alpha_k^2 L\cdot\big\langle \xi_k^{(i)}, \mathsf{x}_k-\mathsf{y}_k\big\rangle + \frac{1}{2}\|\mathsf{z}_{k-1}-u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)}-u\|_A^2 \ .
\end{aligned}$$

Here, ① is due to the minimality of $\mathsf{z}_k^{(i)} = \arg\min_{z\in\Delta_{\mathsf{box}}}\big\{V_{\mathsf{z}_{k-1}}(z) + \big\langle n\alpha_k\xi_k^{(i)}, z\big\rangle\big\}$, which implies that $\big\langle \nabla V_{\mathsf{z}_{k-1}}(\mathsf{z}_k^{(i)}) + n\alpha\xi_k^{(i)}, u-\mathsf{z}_k^{(i)}\big\rangle \ge 0$ for all $u\in\Delta_{\mathsf{box}}$. Step ② is due to the "three-point equality" of Bregman divergence (cf. [13]), which can be checked for every coordinate $\ell\in[n]$ as follows:

$$\begin{aligned}
-\nabla_\ell V_{\mathsf{z}_{k-1}}(\mathsf{z}_k^{(i)})\cdot(\mathsf{z}_{k,\ell}^{(i)}-u_\ell) &= \|A_{:i}\|_\infty(\mathsf{z}_{k-1,\ell}-\mathsf{z}_{k,\ell}^{(i)})\cdot(\mathsf{z}_{k,\ell}^{(i)}-u_\ell) \\
&= \|A_{:i}\|_\infty\Big(-\frac{1}{2}(\mathsf{z}_{k-1,\ell}-\mathsf{z}_{k,\ell}^{(i)})^2 + \frac{1}{2}(u_\ell-\mathsf{z}_{k-1,\ell})^2 - \frac{1}{2}(\mathsf{z}_{k,\ell}^{(i)}-u_\ell)^2\Big) \ .
\end{aligned}$$

③ is by our choice of y_k which satisfies that $\mathsf{z}_{k-1}-\mathsf{z}_k^{(i)} = n\alpha_k L(\mathsf{x}_k-\mathsf{y}_k^{(i)})$. □

Proposition 3.7. *If $\mathsf{z}_{k-1}\in\Delta_{\mathsf{box}}$, the minimizer $z = \arg\min_{z\in\Delta_{\mathsf{box}}}\big\{\frac{1}{2}\|z-\mathsf{z}_{k-1}\|_A^2 + \langle\delta\mathbf{e}_i, z\rangle\big\}$ for any scalar $\delta\in\mathbb{R}$ and basis vector $\mathbf{e}_i$ can be computed as follows:*

1. *$z\leftarrow\mathsf{z}_{k-1}$.*
2. *$z_i\leftarrow z_i-\delta/\|A_{:i}\|_\infty$.*
3. *If $z_i<0$, then $z_i\leftarrow 0$; if $z_i>1/\|A_{:i}\|_\infty$, $z_i\leftarrow 1/\|A_{:i}\|_\infty$.*
4. *Return z.*

Proof of Proposition 3.7. Let us denote by z the returned value of the described procedure, and $g(u)\overset{\text{def}}{=}\frac{1}{2}\|u-\mathsf{z}_{k-1}\|_A^2 + \langle\delta\mathbf{e}_i, u\rangle$. Since Δ_{box} is a convex body and $g(\cdot)$ is convex, to show $z = \arg\min_{z\in\Delta_{\mathsf{box}}}\{g(z)\}$, it suffices for us to prove that for every $u\in\Delta_{\mathsf{box}}$, $\langle\nabla g(z), u-z\rangle\ge 0$. Since the gradient $\nabla g(z)$ can be written explicitly, this is equivalent to

$$\delta(u_i-z_i) + \sum_{\ell=1}^n\|A_{:\ell}\|_\infty\cdot\big(z_\ell-\mathsf{z}_{k-1,\ell}\big)\cdot(u_\ell-z_\ell)\ge 0 \ .$$

However, since $z_\ell = \mathsf{z}_{k-1,\ell}$ for every $\ell\ne i$, this is equivalent to

$$\Big(\delta + \|A_{:i}\|_\infty\cdot\big(z_i-\mathsf{z}_{k-1,i}\big)\Big)\cdot(u_i-z_i)\ge 0 \ .$$

There are three possibilities here. If $z_i = \mathsf{z}_{k-1,i}-\delta/\|A_{:i}\|_\infty$ then the left-hand side is zero and we are done. Otherwise, if $z_i > \mathsf{z}_{k-1,i}-\delta/\|A_{:i}\|_\infty$, then it must satisfy that $z_i = 0$; in such a case the left-hand side is the multiplication of two non-negatives, and therefore non-positive. If $z_i < \mathsf{z}_{k-1,i}-\delta/\|A_{:i}\|_\infty$, then it must satisfy that $z_i = 1/\|A_{:i}\|_\infty$; in such a case the left-hand side is the multiplication of two non-positives, and therefore non-positive. □

Lemma 3.9. *We have $f_\mu(\mathsf{x}_k)-f_\mu(\mathsf{y}_k^{(i)})\ge\frac{1}{2}\langle\nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k-\mathsf{y}_k^{(i)}\rangle$. In particular, this implies $f_\mu(\mathsf{x}_k)\ge f_\mu(\mathsf{y}_k^{(i)})$ because $\nabla_i f_\mu(\mathsf{x}_k)$ and $\mathsf{x}_{k,i}-\mathsf{y}_{k,i}^{(i)}$ have the same sign, while $\mathsf{x}_{k,\ell} = \mathsf{y}_{k,\ell}^{(i)}$ for $\ell\ne i$.*

Proof. Note that $\mathsf{y}_k^{(i)} = \mathsf{x}_k + \lambda \mathbf{e}_i$ for some step length λ such that $|\lambda| \le \frac{1}{L\|A_{:i}\|_\infty}$ according to Fact 3.8. We first prove this lemma in the case of $\nabla_i f_\mu(\mathsf{x}_k) \in [-1,1]$ so that $\xi_{k,i}^{(i)} = \nabla_i f_\mu(\mathsf{x}_k)$.

$$\begin{aligned}
f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) = f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{x}_k + \lambda \mathbf{e}_i) &= -\int_0^\lambda \Big(\nabla_i f_\mu(\mathsf{x}_k + \chi \mathbf{e}_i)\Big) d\chi \\
&\overset{①}{\ge} \int_0^\lambda \Big(-\nabla_i f_\mu(\mathsf{x}_k) - L\|A_{:i}\|_\infty \cdot |\chi|\Big) d\chi \\
&= -\nabla_i f_\mu(\mathsf{x}_k) \cdot |\lambda| - \frac{L\|A_{:i}\|_\infty}{2} \cdot \lambda^2 \\
&\overset{②}{\ge} -\nabla_i f_\mu(\mathsf{x}_k) \cdot |\lambda| - \frac{L\|A_{:i}\|_\infty}{2} \cdot |\lambda| \cdot \frac{|\xi_{k,i}^{(k)}|}{L\|A_{:i}\|_\infty} \\
&= -\frac{1}{2}\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{y}_k^{(i)} - \mathsf{x}_k \rangle \ .
\end{aligned}$$

Above, ① uses Lemma 2.6.a, and ② uses Fact 3.8.

Next, we turn to the case of $\nabla_i f_\mu(\mathsf{x}_k) > 1$.

$$\begin{aligned}
f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) = f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{x}_k + \lambda \mathbf{e}_i) &= -\int_0^\lambda \nabla_i f_\mu(\mathsf{x}_k + \chi \mathbf{e}_i) d\chi \\
&\overset{①}{\ge} \int_0^\lambda \Big(1 - \frac{\|A_{:i}\|_\infty L}{2}|\chi|\Big) \nabla_i f_\mu(x) d\chi \\
&\overset{②}{\ge} \int_0^\lambda \frac{1}{2} \nabla_i f_\mu(x) d\chi = \frac{1}{2}\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \rangle \ .
\end{aligned}$$

Above, ① uses Lemma 2.6.b and ② uses $|\chi| \le |\lambda| \le \frac{1}{L\|A_{:i}\|_\infty}$. □

Lemma 3.10. *For every $u \ge 0$,*

$$\big\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \big\rangle + n^2\alpha_k^2 L \cdot \big\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle \le 3n\alpha_k L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \ .$$

Proof. Now there are three possibilities:

- If $\eta_k^{(i)} = 0$, then we must have $\xi_{k,i}^{(i)} = \nabla_i f_\mu(\mathsf{x}_k) \in [-1,1]$, and Lemma 3.9 immediately implies

$$\begin{aligned}
&\big\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \big\rangle + n^2\alpha_k^2 L \cdot \big\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle \\
&\qquad = n^2\alpha_k^2 L \cdot \big\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle \le 2n^2\alpha_k^2 L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}))
\end{aligned}$$

- If $\eta_k^{(i)} > 0$ and $\mathsf{z}_{k,i}^{(i)} > 0$, then we precisely have $\mathsf{z}_{k,i}^{(i)} = \mathsf{z}_{k-1,i} - \frac{n\alpha_k}{\|A_{:i}\|_\infty}$ (see Proposition 3.7), and accordingly $\mathsf{y}_{k,i}^{(i)} = \mathsf{x}_{k,i} - \frac{1}{L\|A_{:i}\|_\infty} < \mathsf{x}_{k,i}$. In this case,

$$\begin{aligned}
&\big\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \big\rangle + n^2\alpha_k^2 L \cdot \big\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle \\
&\overset{①}{\le} n\alpha_k \cdot \nabla f_\mu(\mathsf{x}_k) \cdot \frac{1}{\|A_{:i}\|_\infty} + n^2\alpha_k^2 L \cdot \big\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle \\
&\overset{②}{<} n\alpha_k \cdot \nabla f_\mu(\mathsf{x}_k) \cdot \frac{1}{\|A_{:i}\|_\infty} + n^2\alpha_k^2 L \cdot \big\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle \\
&\overset{③}{=} n\alpha_k L \cdot \big\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle + n^2\alpha_k^2 L \cdot \big\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle
\end{aligned}$$

$$\overset{④}{\leq} \left(2n\alpha_k L + 2n^2\alpha_k^2 L\right) \cdot \left(f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})\right) \ .$$

Above, ① follows from the fact that $\mathsf{z}_{k-1} \in \Delta_{\mathsf{box}}$ and therefore $\mathsf{z}_{k-1,i} \leq \frac{1}{\|A_{:i}\|_\infty}$ by the definition of Δ_{box}, and $u \geq 0$; ② follows from the fact that x_k and $\mathsf{y}_k^{(i)}$ are only different at coordinate i, and $\xi_{k,i}^{(i)} = 1 < \nabla_i f_\mu(\mathsf{x}_k)$ (since $\eta_{k,i}^{(i)} > 0$); ③ follows from the fact that $\mathsf{y}_k^{(i)} = \mathsf{x}_k - \frac{\mathbf{e}_i}{L\|A_{:i}\|_\infty}$; and ④ uses Lemma 3.9.

- If $\eta_k^{(i)} > 0$ and $\mathsf{z}_{k,i}^{(i)} = 0$, then we have

$$\begin{aligned}
&\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\
&\overset{①}{\leq} \left(n\alpha_k \nabla f_\mu(\mathsf{x}_k) \cdot \mathsf{z}_{k-1,i}\right) + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\
&\overset{②}{=} \left\langle n\alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)} \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\
&\overset{③}{=} n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\
&\overset{④}{\leq} 4n^2\alpha_k^2 L \cdot \left(f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})\right) \ .
\end{aligned}$$

Above, ① is because $u \geq 0$, $\nabla_i f_\mu(\mathsf{x}_k) = \eta_{k,i}^{(i)} + 1 > \eta_{k,i}^{(i)}$ and $\nabla_i f_\mu(\mathsf{x}_k) > \xi_{k,i}^{(i)}$; ② uses the assumption that $\mathsf{z}_{k,i}^{(i)} = 0$ and the fact that $\mathsf{z}_{k-1,\ell} = \mathsf{z}_{k,\ell}^{(i)}$ for every $\ell \neq i$; ③ is from our choice of y_k which satisfies that $\mathsf{z}_{k-1} - \mathsf{z}_k^{(i)} = n\alpha_k L(\mathsf{x}_k - \mathsf{y}_k^{(i)})$; and ④ uses Lemma 3.9.

Combining the three cases above, and using the fact that $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq 0$, we conclude that

$$\begin{aligned}
\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle &\leq (2n\alpha_k L + 4n^2\alpha_k^2 L) \cdot \left(f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})\right) \\
&\leq 3n\alpha_k L \cdot \left(f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})\right) \ .
\end{aligned}$$

Above, the last inequality uses our choice of α_k (see Algorithm 1). □

Corollary 3.4. *With probability at least* 9/10, $\texttt{PacLPSolver}(A, x^{\mathsf{start}}, \varepsilon)$ *outputs a* $(1 - O(\varepsilon))$ *approximate solution to the packing LP program. The expected running time is* $O(\frac{\log(nm/\varepsilon)\log(1/\varepsilon)}{\varepsilon} N)$.

Proof. Since for every $x \in \Delta_{\mathsf{box}}$ it satisfies $f_\mu(x) \geq -(1+\varepsilon)\mathsf{OPT}$ according to Proposition 2.3.b, we obtain that $f_\mu(y_T) + (1+\varepsilon)\mathsf{OPT}$ is a random variable that is non-negative, whose expectation $\mathbf{E}[f_\mu(y_T) + (1+\varepsilon)\mathsf{OPT}] \leq 4\varepsilon$. By Markov bound, with at least probability 9/10, we obtain some y_T satisfying $f_\mu(y_T) \leq -(1 - O(\varepsilon))\mathsf{OPT}$, which yields some $(1 - O(\varepsilon))$ approximate solution according to Proposition 2.3.c.

The running time follows from our efficient implementation in Section G. □

C Missing Proofs for Section 4

Proposition 4.3.

(a) $\mathsf{OPT} \in [1, m]$.
(b) $f_\mu(u^*) \leq (1+\varepsilon)\mathsf{OPT}$ *for* $u^* \overset{\text{def}}{=} (1+\varepsilon/2)x^* \in \Delta_{\mathsf{simplex}}$.
(c) $f_\mu(x) \geq (1-\varepsilon)\mathsf{OPT}$ *for every* $x \geq 0$.
(d) For any $x \geq 0$ *satisfying* $f_\mu(x) \leq 2\mathsf{OPT}$, *we must have* $Ax \geq (1-\varepsilon)\mathbb{1}$.
(e) If $x \geq 0$ *satisfies* $f_\mu(x) \leq (1 + O(\varepsilon))\mathsf{OPT}$, *then* $\frac{1}{1-\varepsilon}x$ *is a* $(1 + O(\varepsilon))$*-approximate solution to the covering LP.*

(f) The gradient of $f_\mu(x)$ can be written as

$$\nabla f_\mu(x) = \mathbb{1} - A^T p(x) \quad \textit{where} \quad p_j(x) \stackrel{\text{def}}{=} e^{\frac{1}{\mu}(1-(Ax)_j)} \tag{C.1}$$

Proof.

(a) Suppose that j^* is the row that achieves the smallest infinite norm $\|A_{j:}\|_\infty$ over all rows $j \in [m]$. Then, for any solution $x \in \mathbb{R}^n_{\geq 0}$ satisfying $\langle A_{:j^*}, x\rangle \geq 1$, we must have $\mathbb{1}^T x \geq 1/\|A_{j^*:}\|_\infty = 1$ using (4.2).

On the other hand, we can construct a feasible solution x as follows. Initialize $x = 0$, and then for each row j, let us find the coordinate i that maximizes the value of A_{ij} among all columns i. Then, we increase x_i by $1/A_{ij} = 1/\|A_{j:}\|_\infty$. After we have exhausted all the m rows, we arrive at some $x \geq 0$ satisfying $Ax \geq \mathbb{1}$ as well as $\mathbb{1}^T x = \sum_j 1/\|A_{j:}\|_\infty \leq m$.

(b) We have $\mathbb{1}^T u^* = (1+\varepsilon/2)\mathsf{OPT}$ by the definition of OPT. Also, from the feasibility constraint $Ax^* \geq \mathbb{1}$ in the covering LP, we have $Au^* - \mathbb{1} \geq \varepsilon/2 \cdot \mathbb{1}$, and can compute $f_\mu(u^*)$ as follows:

$$\begin{aligned} f_\mu(u^*) = \mu \sum_j e^{\frac{1}{\mu}(1-(Au^*)_j)} + \mathbb{1}^T u^* &\leq \mu \sum_j e^{\frac{-\varepsilon/2}{\mu}} + (1+\varepsilon/2)\mathsf{OPT} \\ &\leq \frac{\mu m}{(nm)^2} + (1+\varepsilon/2)\mathsf{OPT} \leq (1+\varepsilon)\mathsf{OPT} \ . \end{aligned}$$

(c) Suppose towards contradiction that $f_\mu(x) < (1-\varepsilon)\mathsf{OPT}$. Since $f_\mu(x) < \mathsf{OPT} \leq m$, we must have that for every $j \in [m]$, it satisfies that $e^{\frac{1}{\mu}(1-(Ax)_j)} \leq f_\mu(x)/\mu \leq m/\mu$. This further implies $(Ax)_j \geq 1-\varepsilon$ by the definition of μ. In other words, $Ax \geq (1-\varepsilon)\mathbb{1}$. By the definition of OPT, we must then have $\mathbb{1}^T x \geq (1-\varepsilon)\mathsf{OPT}$, finishing the proof that $f_\mu(x) \geq \mathbb{1}^T x \geq (1-\varepsilon)\mathsf{OPT}$, giving a contradiction.

(d) To show $Ax \geq (1-\varepsilon)\mathbb{1}$, we can assume that $v = \max_j(1-(Ax)_j) > \varepsilon$ because otherwise we are done. Under this definition, we have

$$f_\mu(x) \geq \mu e^{\frac{v}{\mu}} = \mu\Big((\frac{nm}{\varepsilon})^4\Big)^{v/\varepsilon} \geq \frac{\varepsilon}{4\log(nm)}(\frac{nm}{\varepsilon})^4 \gg 2\mathsf{OPT} \ ,$$

contradicting to our assumption that $f_\mu(x) \leq 2\mathsf{OPT}$. Therefore, we must have $v \leq \varepsilon$, that is, $Ax \geq (1-\varepsilon)\mathbb{1}$.

(e) For any x satisfying $f_\mu(x) \leq (1+O(\varepsilon))\mathsf{OPT} \leq 2\mathsf{OPT}$, owing to Proposition 4.3.d, we first have that x is approximately feasible, i.e., $Ax \geq (1-\varepsilon)\mathbb{1}$. Next, because $\mathbb{1}^T x \leq f_\mu(x) \leq (1+O(\varepsilon))\mathsf{OPT}$, we know that x yields an objective $\mathbb{1}^T x \leq (1+O(\varepsilon))\mathsf{OPT}$. Letting $x' = \frac{1}{1-\varepsilon}x$, we both have that x' is feasible (i.e., $Ax' \geq \mathbb{1}$), and x' has an objective $\mathbb{1}^T x'$ at most $(1+O(\varepsilon))\mathsf{OPT}$.

(f) Straightforward by some simple computation. □

D Missing Proofs for Section 5

In this section we prove Theorem 5.3. Because the proof structure is almost identical to that of Theorem 3.2, we spend most of the discussions only pointing out the difference rather than repeating the proofs.

The following three lemmas are completely identical to the ones in the packing LP case, so we restate them below:

Lemma D.1 (cf. Lemma 3.3)**.** *Each iteration of* $\texttt{CovLPSolver}^{\mathsf{wb}}$ *can be implemented to run in expected* $O(N/n)$ *time.*

Lemma D.2 (cf. Lemma 3.5)**.** *We have* $\mathsf{x}_k, \mathsf{y}_k, \mathsf{z}_k \in \Delta_{\mathsf{box}}$ *for all* $k = 0, 1, \ldots, T$*.*

Lemma D.3 (cf. Lemma 3.6)**.** *For every* $u \in \Delta_{\mathsf{box}}$*, it satisfies* $\left\langle n\alpha_k \xi_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle \le n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle + \frac{1}{2}\|\mathsf{z}_{k-1} - u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)} - u\|_A^2$.

As for the gradient descent guarantee (see Section 3.3), one can prove that Lemma 3.9 remains true, as long as one makes necessary changes to Lemma 2.6 which is used in the proof of Lemma 3.9:

Lemma D.4 (cf. Lemma 2.6)**.** *Define the smoothness parameter* $L \stackrel{\text{def}}{=} \frac{4}{\mu}$*. Then, for every* $x \ge 0$*, every* $i \in [n]$*, and every* $\lambda \ge -\frac{1}{L\|A_{:i}\|_\infty}$*:*

(a) If $|\nabla_i f_\mu(x)| \le 1$*, then* $\left|\nabla_i f_\mu(x + \lambda \mathbf{e}_i) - \nabla_i f_\mu(x)\right| \le L\|A_{:i}\|_\infty \cdot |\lambda|$.

(b) If $\nabla_i f_\mu(x) \le -1$*, then* $\nabla_i f_\mu(x + \lambda \mathbf{e}_i) \le \left(1 - \frac{\|A_{:i}\|_\infty L}{2}|\lambda|\right)\nabla_i f_\mu(x)$.

Lemma D.5 (cf. Lemma 3.9)**.** *We have* $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \ge \frac{1}{2}\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \rangle$*. In particular, this implies* $f_\mu(\mathsf{x}_k) \ge f_\mu(\mathsf{y}_k^{(i)})$ *because* $\nabla_i f_\mu(\mathsf{x}_k)$ *and* $\mathsf{x}_{k,i} - \mathsf{y}_{k,i}^{(i)}$ *have the same sign, while* $\mathsf{x}_{k,\ell} = \mathsf{y}_{k,\ell}^{(i)}$ *for* $\ell \ne i$*.*

Since the proofs of Lemma D.4 and D.5 are completely analogous to those in the packing LP case, we skip the details.

Putting All Together. Denote by $\eta_k^{(i)} \in \mathbb{R}^n_{\le 0}$ the vector that is only non-zero at coordinate i, and satisfies $\eta_{k,i}^{(i)} = \nabla_i f_\mu(\mathsf{x}_k) - \xi_{k,i}^{(i)} \in (-\infty, 0]$. In other words, the full gradient

$$\nabla f_\mu(\mathsf{x}_k) = \mathbf{E}_i[n\nabla_i f_\mu(\mathsf{x}_k)] = \mathbf{E}_i[n\eta_k^{(i)} + n\xi_k^{(i)}]$$

can be (in expectation) decomposed into the a large but non-positive component $\eta_k^{(i)} \in (-\infty, 0]^n$ and a small component $\xi_k^{(i)} \in [-1, 1]^n$.

Similar as Section 3.4, for any $u \in \Delta_{\mathsf{box}}$, we can use a basic convexity argument and the mirror descent lemma to compute that

$$\begin{aligned}
&\alpha_k(f_\mu(\mathsf{x}_k) - f_\mu(u)) \le \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u \right\rangle \\
&= \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{z}_{k-1} \right\rangle + \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - u \right\rangle \\
&= \left\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{z}_{k-1} \right\rangle + \mathbf{E}_i\left[\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + \left\langle n\alpha_k \xi_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle\right] \\
&\stackrel{①}{=} \frac{(1-\tau)\alpha_k}{\tau}\left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{y}_{k-1} - \mathsf{x}_k \right\rangle + \mathbf{E}_i\left[\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + \left\langle n\alpha_k \xi_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle\right] \qquad \text{(D.1)} \\
&\stackrel{②}{\le} \frac{(1-\tau)\alpha_k}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) \\
&\qquad + \mathbf{E}_i\Big[\boxed{\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle} + \frac{1}{2}\|\mathsf{z}_{k-1} - u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k^{(i)} - u\|_A^2\Big] \qquad \text{(D.2)}
\end{aligned}$$

Above, ① is because $\mathsf{x}_k = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$, which implies that $\tau(\mathsf{x}_k - \mathsf{z}_{k-1}) = (1-\tau)(\mathsf{y}_{k-1} - \mathsf{x}_k)$. ② uses convexity and Lemma D.3.

We can establish the following lemma to upper bound the boxed term in (D.2). Its proof is in the same spirit to that of Lemma 3.10, and is the *only* place that we require all vectors to reside in Δ_{box}.

Lemma D.6 (cf. Lemma 3.10)**.** *For every $u \in \Delta_{\mathsf{box}}$,*

$$\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \le 21 n\alpha_k L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \ .$$

Proof of Lemma D.6. Now there are three possibilities:

- If $\eta_k^{(i)} = 0$, then we must have $\xi_{k,i}^{(i)} = \nabla_i f_\mu(\mathsf{x}_k) \in [-1, 1]$, and Lemma D.5 immediately implies

$$\begin{aligned}&\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\qquad = n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \le 2n^2\alpha_k^2 L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}))\end{aligned}$$

- If $\eta_k^{(i)} < 0$ and $\mathsf{z}_{k,i}^{(i)} < \frac{10}{\|A_{:i}\|_\infty}$ (and thus $\mathsf{z}_k^{(i)}$ is not on the boundary of Δ_{box}), then we precisely have $\mathsf{z}_{k,i}^{(i)} = \mathsf{z}_{k-1,i} + \frac{n\alpha_k}{\|A_{:i}\|_\infty}$, and accordingly $\mathsf{y}_{k,i}^{(i)} = \mathsf{x}_{k,i} + \frac{1}{L\|A_{:i}\|_\infty} > \mathsf{x}_{k,i}$. In this case,

$$\begin{aligned}&\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{①}{\le} n\alpha_k \cdot \nabla f_\mu(\mathsf{x}_k) \cdot \frac{-10}{\|A_{:i}\|_\infty} + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{②}{<} n\alpha_k \cdot \nabla f_\mu(\mathsf{x}_k) \cdot \frac{-10}{\|A_{:i}\|_\infty} + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{③}{=} 10n\alpha_k L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{④}{\le} \left(20n\alpha_k L + 2n^2\alpha_k^2 L\right) \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \ .\end{aligned}$$

  Above, ① follows from the fact that $\mathsf{z}_{k-1} \in \Delta_{\mathsf{box}}$ and therefore $\mathsf{z}_{k-1,i} \le \frac{10}{\|A_{:i}\|_\infty}$ by the definition of Δ_{box}, and $u \ge 0$; ② follows from the fact that x_k and $\mathsf{y}_k^{(i)}$ are only different at coordinate i, and $\xi_{k,i}^{(i)} = -1 > \nabla_i f_\mu(\mathsf{x}_k)$ (since $\eta_{k,i}^{(i)} < 0$); ③ follows from the fact that $\mathsf{y}_k^{(i)} = \mathsf{x}_k + \frac{\mathbf{e}_i}{L\|A_{:i}\|_\infty}$; and ④ uses Lemma D.5.

- If $\eta_k^{(i)} < 0$ and $\mathsf{z}_{k,i}^{(i)} = \frac{10}{\|A_{:i}\|_\infty}$, then we have

$$\begin{aligned}&\left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{①}{\le} \left\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)} \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{②}{\le} \left\langle n\alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)} \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{③}{=} n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle + n^2\alpha_k^2 L \cdot \left\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{y}_k^{(i)} \right\rangle \\ &\overset{④}{\le} 4n^2\alpha_k^2 L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \ .\end{aligned}$$

  Above, ① is because $u_i \le \frac{10}{\|A_{:i}\|_\infty} = \mathsf{z}_{k,i}^{(i)}$ and $\eta_{k,i}^{(i)} < 0$, together with $\nabla_i f_\mu(\mathsf{x}_k) < \xi_{k,i}^{(i)}$ and $\mathsf{x}_{k,i} \le \mathsf{y}_{k,i}^{(i)}$; ② uses $\nabla_i f_\mu(\mathsf{x}_k) = \eta_{k,i}^{(i)} - 1 < \eta_{k,i}^{(i)}$ and $\mathsf{z}_{k,i}^{(i)} \ge \mathsf{z}_{k-1,i}$; ② uses the assumption that $\mathsf{z}_{k,i}^{(i)} = 0$ and the fact that $\mathsf{z}_{k-1,\ell} = \mathsf{z}_{k,\ell}^{(i)}$ for every $\ell \ne i$; ③ is from our choice of y_k which satisfies that $\mathsf{z}_{k-1} - \mathsf{z}_k^{(i)} = n\alpha_k L(\mathsf{x}_k - \mathsf{y}_k^{(i)})$; and ④ uses Lemma D.5.

Combining the three cases above, and using the fact that $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq 0$, we conclude that

$$\begin{aligned}\big\langle n\alpha_k \eta_k^{(i)}, \mathsf{z}_{k-1} - u \big\rangle + n^2\alpha_k^2 L \cdot \big\langle \xi_k^{(i)}, \mathsf{x}_k - \mathsf{y}_k^{(i)} \big\rangle &\leq (20n\alpha_k L + 4n^2\alpha_k^2 L) \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \\ &\leq 21n\alpha_k L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) \ .\end{aligned}$$

Above, the last inequality uses our choice of α_k (see Algorithm 2). □

Plugging the above lemma into (D.2), we have

$$\begin{aligned}&\alpha_k(f_\mu(\mathsf{x}_k) - f_\mu(u)) \leq \big\langle \alpha_k \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u \big\rangle \\ &\overset{①}{\leq} \frac{(1-\tau)\alpha_k}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \mathbf{E}_i\Big[21n\alpha_k L \cdot (f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)})) + \frac{1}{2}\|\mathsf{z}_{k-1} - u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k - u\|_A^2\Big] \\ &\overset{②}{\leq} \alpha_k f_\mu(\mathsf{x}_k) + \big(21n\alpha_k L - \alpha_k\big) f_\mu(\mathsf{y}_{k-1}) + \mathbf{E}_i\Big[- 21n\alpha_k L \cdot f_\mu(\mathsf{y}_k^{(i)}) + \frac{1}{2}\|\mathsf{z}_{k-1} - u\|_A^2 - \frac{1}{2}\|\mathsf{z}_k - u\|_A^2\Big] \ .\end{aligned} \tag{D.3}$$

Above, ① is because we have chosen α_k so that $n\alpha_k \leq n\alpha_T = \frac{1}{\varepsilon L} \leq \frac{1}{4}$; and ② is because we have chosen τ to satisfy $\frac{1}{\tau} = 21nL$.

Next, recall that we have picked α_k so that $(21nL - 1)\alpha_k = 21nL \cdot \alpha_{k-1}$ in $\texttt{CovLPSolver}^{\mathsf{wb}}$. Telescoping (D.3) for $k = 1, \ldots, T$ and choosing $u^* = (1 + \varepsilon/2)x^*$, we have

$$\begin{aligned}-\sum_{k=1}^{T} \alpha_k f_\mu(u^*) &\leq 21 f_\mu(\mathsf{y}_0) - 21n\alpha_T L \cdot \mathbf{E}[f_\mu(\mathsf{y}_T)] + \|\mathsf{z}_0 - u^*\|_A^2 \\ &\leq -21n\alpha_T L \cdot \mathbf{E}[f_\mu(\mathsf{y}_T)] + 75\mathsf{OPT} \ .\end{aligned}$$

Here, the second inequality is due to $f_\mu(\mathsf{y}_0) = f_\mu(x^{\mathsf{start}}) \leq 3\mathsf{OPT}$ from Fact 5.2, and the fact that

$$\begin{aligned}\|\mathsf{z}_0 - u^*\|_A^2 = \|u^*\|_A^2 = \sum_{i=1}^{n} (u_i^*)^2 \cdot \|A_{:i}\|_\infty &\leq (1+\varepsilon/2)^2 \sum_{i=1}^{n} (x_i^*)^2 \cdot \|A_{:i}\|_\infty \\ &\leq 10(1+\varepsilon/2)^2 \sum_{i=1}^{n} x_i^* < 12\mathsf{OPT} \ .\end{aligned}$$

Finally, using the fact that $\sum_{k=1}^{T} \alpha_k = \alpha_T \cdot \sum_{k=0}^{T-1} \big(1 - \frac{1}{21nL}\big)^k = 21n\alpha_T L\big(1 - (1 - \frac{1}{21nL})^T\big)$, we rearrange and obtain that

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \leq \frac{\sum_k \alpha_k}{21n\alpha_T L} f_\mu(u^*) + \frac{75}{21n\alpha_T L}\mathsf{OPT} = \big(1 - (1 - \frac{1}{21nL})^T\big) f_\mu(u^*) + \frac{75}{21n\alpha_T L}\mathsf{OPT} \ .$$

Choosing $T = \lceil 21nL \log(1/\varepsilon) \rceil$ so that $\frac{1}{n\alpha_T L} = (1 - \frac{1}{21nL})^T \leq \varepsilon$. Combining this with the fact that $f_\mu(u^*) \leq (1+\varepsilon)\mathsf{OPT}$ (see Proposition 4.3.b), we obtain

$$\mathbf{E}[f_\mu(\mathsf{y}_T)] \leq (1+\varepsilon)\mathsf{OPT} + 3.6\varepsilon \cdot \mathsf{OPT} < (1 + 4.6\varepsilon)\mathsf{OPT} \ .$$

Therefore, we have finished proving Theorem 5.3. □

E Missing Proofs for Section 6

Proposition 6.1. *Letting* $x^{\mathsf{start}} = (1+\varepsilon/2)\cdot x^\sharp + (\frac{1}{n},\dots,\frac{1}{n})$, *we have* $\mathbb{1}^T x^{\mathsf{start}} \le 2\mathsf{OPT}'$ *and* $f_\mu(x^{\mathsf{start}}) \le 4\mathsf{OPT}$.

Proof. Using the fact that $Ax^{\mathsf{start}} - \mathbb{1} \ge (1+\varepsilon/2)Ax^\sharp - \mathbb{1} \ge \varepsilon/2\cdot\mathbb{1}$, we compute $f_\mu(x^{\mathsf{start}})$ as follows:

$$\begin{aligned} f_\mu(x^{\mathsf{start}}) &= \mu\sum_j e^{\frac{1}{\mu}(1-(Ax^{\mathsf{start}})_j)} + \mathbb{1}^T x^{\mathsf{start}} \le \mu\sum_j e^{\frac{-\varepsilon/2}{\mu}} + 2\mathsf{OPT} + 1 \\ &\le \frac{\mu m}{(nm)^2} + 3\mathsf{OPT} < 4\mathsf{OPT}\ . \end{aligned}$$

Also, we have $\mathbb{1}^T x^{\mathsf{start}} \le (1+\varepsilon/2)\mathsf{OPT}' + 1 \le 2\mathsf{OPT}'$. □

Lemma 6.8.

$$\begin{aligned} f_\mu(\mathsf{x}_k) - f_\mu(u^*) &\le \langle \mathbb{1} - A^T p(\mathsf{x}_k), \mathsf{x}_k - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle + \varepsilon\mathsf{OPT} \\ &= \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle + \varepsilon\mathsf{OPT} \end{aligned}$$

Proof.

$$\begin{aligned} f_\mu(\mathsf{x}_k) - f_\mu(u^*) &= \mu\sum_{j=1}^m \left(e^{\frac{1}{\mu}(1-(A\mathsf{x}_k)_j)} - e^{\frac{1}{\mu}(1-(Au^*)_j)}\right) + \langle \mathbb{1}, \mathsf{x}_k - u^*\rangle \\ &\overset{①}{\le} \mu\sum_{j=1}^m \left(e^{\frac{1}{\mu}(1-(A\mathsf{x}_k)_j)} - e^{\frac{1}{\mu}(1-(\widetilde{A}u^*)_j)}\right) + \langle \mathbb{1}, \mathsf{x}_k - u^*\rangle + \mu\cdot m\cdot e^{-1/\mu} \\ &\overset{②}{\le} \sum_{j=1}^m e^{\frac{1}{\mu}(1-(A\mathsf{x}_k)_j)}\cdot\big((\widetilde{A}u^*)_j - (A\mathsf{x}_k)_j\big) + \langle \mathbb{1}, \mathsf{x}_k - u^*\rangle + \varepsilon\mathsf{OPT} \\ &= \sum_{j=1}^m p_j(\mathsf{x}_k)\cdot\big((\widetilde{A}u^*)_j - (A\mathsf{x}_k)_j\big) + \langle \mathbb{1}, \mathsf{x}_k - u^*\rangle + \varepsilon\mathsf{OPT} \\ &= \sum_{j=1}^m p_j(\mathsf{x}_k)\cdot\big((Au^*)_j - (A\mathsf{x}_k)_j\big) + \langle \mathbb{1}, \mathsf{x}_k - u^*\rangle + \sum_{j=1}^m p_j(\mathsf{x}_k)\cdot\big((\widetilde{A}u^*)_j - (Au^*)_j\big) + \varepsilon\mathsf{OPT} \\ &= \langle -Ap(\mathsf{x}_k), \mathsf{x}_k - u^*\rangle + \langle \mathbb{1}, \mathsf{x}_k - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle + \varepsilon\mathsf{OPT}\ . \end{aligned}$$

Above, ① is because if $(Au^*)_j \ne (\widetilde{A}u^*)_j$ for some j, then it must satisfy that $(\widetilde{A}u^*)_j = 2$, and therefore $-e^{\frac{1}{\mu}(1-(Au^*)_j)} \le -e^{\frac{1}{\mu}(1-(\widetilde{A}u^*)_j)} + e^{-1/\mu}$. ② uses the convexity inequality of $e^b - e^a \le \langle e^b, b-a\rangle$, and the fact that $\mu m e^{-1/\mu} \ll \varepsilon\mathsf{OPT}$. □

Lemma 6.10.

$$f_\mu(\mathsf{x}_k) - f_\mu(u^*) \le \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \mathbf{E}_i\Big[\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\rangle\Big] + \mathbf{E}_i\Big[\langle n\widetilde{\eta}_k^{(i)}, -u^*\rangle\Big] + \varepsilon\mathsf{OPT}\ .$$

Proof.

$$\begin{aligned} &\big(f_\mu(\mathsf{x}_k) - f_\mu(u^*)\big) - \varepsilon\mathsf{OPT} \\ &\overset{①}{\le} \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle \end{aligned}$$

$$
\begin{aligned}
&= \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{x}_k - \mathsf{z}_{k-1}\rangle + \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle \\
&\overset{②}{=} \frac{(1-\tau)}{\tau}\langle \nabla f_\mu(\mathsf{x}_k), \mathsf{y}_{k-1} - \mathsf{x}_k\rangle + \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle \\
&\overset{③}{\le} \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \langle \nabla f_\mu(\mathsf{x}_k), \mathsf{z}_{k-1} - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle \\
&= \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \langle \xi_k + \eta_k, \mathsf{z}_{k-1} - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k), u^*\rangle \\
&\overset{④}{\le} \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \langle \xi_k, \mathsf{z}_{k-1} - u^*\rangle + \langle \widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k) - \eta_k, u^*\rangle \\
&\overset{⑤}{\le} \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \langle \xi_k, \mathsf{z}_{k-1} - u^*\rangle + \langle -\widetilde{\eta}_k, u^*\rangle \\
&= \frac{(1-\tau)}{\tau}(f_\mu(\mathsf{y}_{k-1}) - f_\mu(\mathsf{x}_k)) + \mathbf{E}_i\Big[\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\rangle + \langle -n\widetilde{\eta}_k^{(i)}, u^*\rangle\Big] \ .
\end{aligned}
$$

Above, ① is due to Lemma 6.8. ② is because $\mathsf{x}_k = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$, which implies that $\tau(\mathsf{x}_k - \mathsf{z}_{k-1}) = (1-\tau)(\mathsf{y}_{k-1} - \mathsf{x}_k)$. ③ is by the convexity of $f_\mu(\cdot)$. ④ is because $\langle \eta_k, \mathsf{z}_{k-1}\rangle \le 0$, since $\eta_k \le 0$ while $\mathsf{z}_{k-1} \ge 0$.

⑤ needs some careful justification: for every $i \notin B_k$, we have $(\widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k))_i - \eta_{k,i} \le 0 - 0 = -\widetilde{\eta}_{k,i}$; for every $i \in B_k$, we have

$$
\begin{aligned}
(\widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k))_i - \eta_{k,i} &= (\widetilde{A}^T p(\mathsf{x}_k) - A^T p(\mathsf{x}_k))_i - \big((1+\beta) - (A^T p(\mathsf{x}_k))_i\big) \\
&= -\big((1+\beta) - (\widetilde{A}^T p(\mathsf{x}_k))_i\big) = -\widetilde{\eta}_{k,i} \ ,
\end{aligned}
$$

where the two equalities follow from the definitions of $\eta_{k,i}$ and $\widetilde{\eta}_{k,i}$ (see Definition 6.9). □

Lemma 6.11. *Denoting by* $\gamma \overset{\text{def}}{=} 2\alpha_T n$, *we have*

$$
\mathbf{E}_i\big[\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\big\rangle\big] \le V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - \mathbf{E}_i\Big[V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big)\Big] + 12\mathsf{OPT}\cdot\gamma\alpha_k\beta \ .
$$

Proof. Define $w(x) \overset{\text{def}}{=} \sum_i x_i\log(x_i) - x_i$ and accordingly, $V_x(y) = w(y) - \langle w'(x), y - x\rangle - w(x) = \sum_i y_i \log\frac{y_i}{x_i} + x_i - y_i$. We first compute using the classical analysis of mirror descent step as follows:

$$
\begin{aligned}
&\gamma\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1}\big\rangle + \alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^*\big\rangle \\
&= (1+\gamma)\alpha_k\Big\langle n\xi_k^{(i)}, \mathsf{z}_k^{(i)} - \frac{u^*}{1+\gamma}\Big\rangle + (1+\gamma)\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)}\big\rangle \\
&\overset{①}{\le} \Big\langle w'(\mathsf{z}_{k-1}) - w'(\mathsf{z}_k^{(i)}), \mathsf{z}_k^{(i)} - \frac{u^*}{1+\gamma}\Big\rangle + (1+\gamma)\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)}\big\rangle \\
&= \Big(w\big(\frac{u^*}{1+\gamma}\big) - w(\mathsf{z}_{k-1}) - \Big\langle w'(\mathsf{z}_{k-1}), \frac{u^*}{1+\gamma} - \mathsf{z}_{k-1}\Big\rangle\Big) - \Big(w\big(\frac{u^*}{1+\gamma}\big) - w(\mathsf{z}_k^{(i)}) - \Big\langle w'(\mathsf{z}_k^{(i)}), \frac{u^*}{1+\gamma} - \mathsf{z}_k^{(i)}\Big\rangle\Big) \\
&\quad + \Big(w(\mathsf{z}_{k-1}) - w(\mathsf{z}_k^{(i)}) - \big\langle w'(\mathsf{z}_{k-1}), \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)}\big\rangle\Big) + (1+\gamma)\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)}\big\rangle \\
&= V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big) + \boxed{(1+\gamma)\alpha_k\big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)}\big\rangle - V_{\mathsf{z}_{k-1}}(\mathsf{z}_k^{(i)})} \ .
\end{aligned} \tag{E.1}
$$

Above, ① is because $\mathsf{z}_k^{(i)} = \arg\min_{z\in\Delta_{\mathsf{simplex}}}\big\{V_{\mathsf{z}_{k-1}}(z) + \langle(1+\gamma)\alpha_k n\xi_k^{(i)}, z\rangle\big\}$, which is equivalent to saying

$$
\forall u \in \Delta_{\mathsf{simplex}}, \quad \langle V'_{\mathsf{z}_{k-1}}(\mathsf{z}_k^{(i)}) + (1+\gamma)\alpha_k n\xi_k^{(i)}, u - \mathsf{z}_k^{(i)}\rangle \ge 0
$$

$$\iff \quad \forall u \in \Delta_{\mathsf{simplex}}, \quad \langle w'(\mathsf{z}_k^{(i)}) - w'(\mathsf{z}_{k-1}) + (1+\gamma)\alpha_k n \xi_k^{(i)}, u - \mathsf{z}_k^{(i)} \rangle \geq 0 \ .$$

In particular, we have $\mathbb{1}^T \frac{u^*}{1+\gamma} = \mathbb{1}^T \frac{(1+\varepsilon/2)x^*}{1+\gamma} < 2\mathsf{OPT} \leq 2\mathsf{OPT}'$ and therefore substituting $u = \frac{u^*}{1+\gamma} \in \Delta_{\mathsf{simplex}}$ into the above inequality we get ①.

Next, we upper bound the term in the box:

$$\begin{aligned}
&(1+\gamma)\alpha_k \langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - \mathsf{z}_k^{(i)} \rangle - V_{\mathsf{z}_{k-1}}(\mathsf{z}_k^{(i)}) \\
&\overset{①}{\leq} (1+\gamma)\alpha_k n \xi_{k,i} \cdot (\mathsf{z}_{k-1,i} - \mathsf{z}_{k,i}^{(i)}) - \Big(\mathsf{z}_{k,i}^{(i)} \log \frac{\mathsf{z}_{k,i}^{(i)}}{\mathsf{z}_{k-1,i}} + \mathsf{z}_{k-1,i} - \mathsf{z}_{k,i}^{(i)} \Big) \\
&\overset{②}{\leq} (1+\gamma)\alpha_k n \xi_{k,i} \cdot (\mathsf{z}_{k-1,i} - \mathsf{z}_{k,i}^{(i)}) - \frac{|\mathsf{z}_{k,i}^{(i)} - \mathsf{z}_{k-1,i}|^2}{2\max\{\mathsf{z}_{k,i}^{(i)}, \mathsf{z}_{k-1,i}\}} \\
&\overset{③}{\leq} (1+\gamma)\alpha_k n \xi_{k,i} \cdot (\mathsf{z}_{k-1,i} - \mathsf{z}_{k,i}^{(i)}) - \frac{|\mathsf{z}_{k,i}^{(i)} - \mathsf{z}_{k-1,i}|^2}{4\mathsf{z}_{k-1,i}} \\
&\overset{④}{\leq} (1+\gamma)^2 \mathsf{z}_{k-1,i} \cdot (\alpha_k n \xi_{k,i})^2 \overset{⑤}{\leq} 2\mathsf{z}_{k-1,i} \cdot (\alpha_k n \xi_{k,i})^2 \overset{⑥}{\leq} \mathsf{z}_{k-1,i} \cdot \gamma\alpha_k n |\xi_{k,i}| \\
&\overset{⑦}{\leq} \mathsf{z}_{k-1,i} \cdot \gamma\alpha_k n \xi_{k,i} + 2\mathsf{z}_{k-1,i} \cdot \gamma\alpha_k n\beta = \gamma\alpha_k \langle n\xi_k^{(i)}, \mathsf{z}_{k-1} \rangle + 2\mathsf{z}_{k-1,i} \cdot \gamma\alpha_k n\beta \ .
\end{aligned} \tag{E.2}$$

Above, ① uses the fact that for every $i' \neq i$, $\mathsf{z}_{k,i'}^{(i)} \log \frac{\mathsf{z}_{k,i'}^{(i)}}{\mathsf{z}_{k-1,i'}} + \mathsf{z}_{k-1,i'} - \mathsf{z}_{k,i}^{(i)} \geq 0$. ② uses the inequality that for every $a, b > 0$, we have $a \log \frac{a}{b} + b - a \geq \frac{(a-b)^2}{2\max\{a,b\}}$.[15] ③ uses the fact that $\mathsf{z}_{k,i}^{(i)} \leq 2\mathsf{z}_{k-1,i}$.[16] ④ uses Cauchy-Shwarz: $ab - b^2/4 \leq a^2$. ⑤ uses $(1+\gamma)^2 < 2$. ⑥ uses $|\xi_{k,i}| \leq 1$ and $\gamma = 2\alpha_T n \geq 2\alpha_k n$. ⑦ uses $\xi_{k,i} \geq -\beta$.

Next, we combine (E.1) and (E.2) to conclude that

$$\alpha_k \big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^* \big\rangle \leq V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big) + 2\mathsf{z}_{k-1,i} \cdot \gamma\alpha_k n\beta \ .$$

Taking expectation on both sides with respect to i, and using the property that $\mathbb{1}^T \mathsf{z}_{k-1} \leq 3\mathsf{OPT}' \leq 6\mathsf{OPT}$, we obtain that

$$\mathbf{E}_i\big[\alpha_k \big\langle n\xi_k^{(i)}, \mathsf{z}_{k-1} - u^* \big\rangle\big] \leq V_{\mathsf{z}_{k-1}}\big(\frac{u^*}{1+\gamma}\big) - \mathbf{E}_i\Big[V_{\mathsf{z}_k^{(i)}}\big(\frac{u^*}{1+\gamma}\big)\Big] + 12\mathsf{OPT} \cdot \gamma\alpha_k\beta \ . \qquad \square$$

Proposition 6.12. *If $\mathsf{z}_{k-1} \in \Delta_{\mathsf{simplex}}$ and $\mathsf{z}_{k-1} > 0$, the minimizer $z = \arg\min_{z \in \Delta_{\mathsf{simplex}}} \big\{ V_{\mathsf{z}_{k-1}}(z) + \langle \delta \mathbf{e}_i, z \rangle \big\}$ for any scalar $\delta \in \mathbb{R}$ and basis vector $\mathbf{e}_i$ can be computed as follows:*

1. *$z \leftarrow \mathsf{z}_{k-1}$.*
2. *$z_i \leftarrow z_i \cdot e^{-\delta}$.*
3. *If $\mathbb{1}^T z > 2\mathsf{OPT}'$, $z \leftarrow \frac{2\mathsf{OPT}'}{\mathbb{1}^T z} z$.*
4. *Return z.*

Proof. Let us denote by z the returned value of the described procedure, and $g(u) \overset{\text{def}}{=} V_{\mathsf{z}_{k-1}}(u) + \langle \delta \mathbf{e}_i, u \rangle$. Since $\Delta_{\mathsf{simplex}}$ is a convex body and $g(\cdot)$ is convex, to show $z = \arg\min_{z \in \Delta_{\mathsf{simplex}}} \{g(u)\}$, it

[15] This inequality in fact corresponds to a local strong convexity property of $w(\cdot)$. We have used this technique in our paper [4].

[16] This is because, our parameter choices ensure that $(1+\gamma)\alpha_k n < 1/2\beta$, which further means $-(1+\gamma)\alpha_k n \xi_k^{(i)} \leq 1/2$. As a result, we must have $\mathsf{z}_{k,i}^{(i)} \leq \mathsf{z}_{k-1,i} \cdot e^{0.5} < 2\mathsf{z}_{k-1,i}$ (see the explicit definition of the mirror step at Proposition 6.12).

suffices for us to prove that for every $u \in \Delta_{\mathsf{simplex}}$, $\langle \nabla g(z), u - z \rangle \geq 0$. Since the gradient $\nabla g(z)$ can be written explicitly, this is equivalent to

$$\delta(u_i - z_i) + \sum_{\ell=1}^{n} \log \frac{z_\ell}{\mathsf{z}_{k-1,\ell}} \cdot (u_\ell - z_\ell) \geq 0 \ .$$

If the re-scaling in step 3 is not executed, then we have $z_\ell = \mathsf{z}_{k-1,\ell}$ for every $\ell \neq i$, and $z_i = \mathsf{z}_{k-1,i} \cdot e^{-\delta}$; thus, the left-hand side is zero so the above inequality is true for every $u \in \Delta_{\mathsf{simplex}}$.

Otherwise, we have $\mathbb{1}^T z = 2\mathsf{OPT}'$ and there exists some constant factor $Z > 1$ such that, $z_\ell = \mathsf{z}_{k-1,\ell}/Z$ for every $\ell \neq i$, and $z_i = \mathsf{z}_{k-1,i} \cdot e^{-\delta}/Z$. In such a case, the left-hand side equals to

$$(u_i - z_i) \cdot (\delta - \delta) + \sum_{\ell=1}^{n} -\log Z \cdot (u_\ell - z_\ell) \ .$$

It is clear at this moment that since $\log Z > 0$ and $\mathbb{1}^T u \leq 2\mathsf{OPT}' = \mathbb{1}^T z$, the above quantity is always non-negative, finishing the proof. □

Lemma 6.13. *For every $i \in [n]$, we have*

(a) $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq 0$, *and*

(b) $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq \frac{\mu\beta}{12} \cdot \langle -\widetilde{\eta}_k^{(i)}, u^* \rangle$.

Proof of Lemma 6.13 part (a). Since if $i \notin B_k$ is not a large index we have $\mathsf{y}_k^{(i)} = \mathsf{x}_k$ and the claim is trivial, we focus on $i \in B_k$ in the remaining proof. Recall that $\mathsf{y}_k^{(i)} = \mathsf{x}_k + \delta \mathbf{e}_i$ for some $\delta > 0$ defined in Algorithm 3, so we have

$$f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) = \int_{\tau=0}^{\delta} \langle -\nabla f_\mu(\mathsf{x}_k + \tau \mathbf{e}_i), \mathbf{e}_i \rangle d\tau = \int_{\tau=0}^{\delta} \big(\langle A_{:i}, p(\mathsf{x}_k + \tau \mathbf{e}_i) \rangle - 1 \big) d\tau \ .$$

It is clear that $\langle A_{:i}, p(\mathsf{x}_k + \tau \mathbf{e}_i) \rangle$ decreases as τ increases, and therefore it suffices to prove that $\langle A_{:i}, p(\mathsf{x}_k + \delta \mathbf{e}_i) \rangle \geq 1$.

Suppose that the rows of $A_{:i}$ are sorted (for the simplicity of notation) by the increasing order of $A_{j,i}$. Now, by the definition of the algorithm, there exists some $j^* \in [m]$ satisfying that

$$\sum_{j<j^*} A_{j,i} \cdot p_j(\mathsf{x}_k) < 1 + \beta \quad \text{and} \quad \sum_{j \leq j^*} A_{j,i} \cdot p_j(\mathsf{x}_k) \geq 1 + \beta \ .$$

Next, by our choice of δ which satisfies $\delta = \frac{\mu\beta}{2A_{j^*,i}} \leq \frac{\mu\beta}{2A_{j,i}}$ for every $j \leq j^*$, we have

$$p_j(\mathsf{x}_k + \delta \mathbf{e}_i) = p_j(\mathsf{x}_k) \cdot e^{-\frac{A_{j,i}\delta}{\mu}} \geq p_j(\mathsf{x}_k) \cdot e^{-\beta/2} \geq p_j(\mathsf{x}_k) \cdot (1 - \beta/2) \ ,$$

and as a result,

$$\langle A_{:i}, p(\mathsf{x}_k + \delta \mathbf{e}_i) \geq \sum_{j \leq j^*} A_{j,i} \cdot p_j(\mathsf{x}_k + \delta \mathbf{e}_i) \geq (1 - \beta/2) \sum_{j \leq j^*} A_{j,i} \cdot p_j(\mathsf{x}_k) \geq (1 - \beta/2)(1 + \beta) \geq 1 \ . \ \square$$

Proof of Lemma 6.13 part (b). Owing to part (a), for every coordinate i such that $\widetilde{\eta}_{k,i} \geq 0$, we automatically have $f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \geq 0$ so the lemma is obvious. Therefore, let us focus only on

coordinates i such that $\widetilde{\eta}_{k,i} < 0$; these are necessarily large indices $i \in B$. Recall from Definition 6.9 that $\widetilde{\eta}_{k,i} = (1+\beta) - (\widetilde{A}^T p(\mathsf{x}_k))_i$, so we have

$$\sum_{j=1}^{m} \widetilde{A}_{j,i} \cdot p_j(\mathsf{x}_k) - (1+\beta) > 0 \ .$$

For the simplicity of description, suppose again that the rows of the i-th column is sorted in the non-decreasing order of $A_{j,i}$. That is, $A_{1,i} \leq \cdots A_{m,i}$. The definition of j^* can be simplified as

$$\sum_{j<j^*} A_{j,i} \cdot p_j(\mathsf{x}_k) < 1+\beta \quad \text{and} \quad \sum_{j\leq j^*} A_{j,i} \cdot p_j(\mathsf{x}_k) \geq 1+\beta \ .$$

Let $j^\flat \in [m]$ be the row such that

$$\sum_{j<j^\flat} \widetilde{A}_{j,i} \cdot p_j(\mathsf{x}_k) < 1+\beta \quad \text{and} \quad \sum_{j\leq j^\flat} \widetilde{A}_{j,i} \cdot p_j(\mathsf{x}_k) \geq 1+\beta \ .$$

Note that such a $j^\flat$ must exist because $\sum_{j=1}^m \widetilde{A}_{j,i} \cdot p_j > 1+\beta$. It is clear that $j^\flat \geq j^*$, owing to the definition that $\widetilde{A}_{ji} \leq A_{ji}$ for all $i \in [n], j \in [m]$. Defining $\delta^\flat = \frac{\mu\beta}{2A_{j^\flat,i}} \leq \delta$, the objective decrease is lower bounded as

$$\begin{aligned}
f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) &= \int_{\tau=0}^{\delta} \langle -\nabla f_\mu(\mathsf{x}_k + \tau\mathbf{e}_i), \mathbf{e}_i\rangle d\tau = \int_{\tau=0}^{\delta} \big(\langle A_{:i}, p(\mathsf{x}_k+\tau\mathbf{e}_i)\rangle - 1\big) d\tau \\
&\geq \int_{\tau=0}^{\delta^\flat} \big(\langle A_{:i}, p(\mathsf{x}_k+\tau\mathbf{e}_i)\rangle - 1\big) d\tau \\
&= \underbrace{\int_{\tau=0}^{\delta^\flat} \Big(\sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k+\tau\mathbf{e}_i) - 1\Big) d\tau}_{I} + \underbrace{\sum_{j>j^\flat} \int_{\tau=0}^{\delta^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k+\tau\mathbf{e}_i) d\tau}_{I'}
\end{aligned}$$

where the inequality is because $\delta^\flat \leq \delta$ and $\langle A_{:i}, p(\mathsf{x}_k+\tau\mathbf{e}_i)\rangle \geq 1$ for all $\tau \leq \delta$ (see the proof of part (a)).

Part *I*. To lower bound I, we use the monotonicity of $p_j(\cdot)$ and obtain that

$$I = \int_{\tau=0}^{\delta^\flat} \Big(\sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k+\tau\mathbf{e}_i) - 1\Big) d\tau \geq \delta^\flat \cdot \Big(\sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k+\delta^\flat\mathbf{e}_i) - 1\Big) \ .$$

However, our choice of $\delta^\flat = \frac{\mu\beta}{2A_{j^\flat,i}} \leq \frac{\mu\beta}{2A_{j,i}}$ for all $j \leq j^\flat$ ensures that

$$\sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k+\delta^\flat\mathbf{e}_i) \geq \sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k)\cdot e^{\frac{-A_{j,i}\cdot\delta^\flat}{\mu}} \geq \sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k)\cdot(1-\beta/2) \ .$$

Therefore, we obtain that

$$I \geq \delta^\flat\Big((1-\beta/2)\sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k) - 1\Big) \geq \frac{\delta^\flat}{3}\Big(\sum_{j\leq j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k) - 1\Big) \ ,$$

where the inequality is because $\left(\frac{2}{3}-\frac{\beta}{2}\right)\sum_{j\le j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k) \ge \frac{4-3\beta}{6}\cdot(1+\beta) \ge \frac{2}{3}$ whenever $\beta \le \frac{1}{3}$ (or equivalently, whenever $\varepsilon \le 1/9$).

Now, suppose that $\sum_{j\le j^\flat} \widetilde{A}_{j,i}\cdot p_j(\mathsf{x}_k) - (1+\beta) = b\cdot \widetilde{A}_{j^\flat,i}\cdot p_{j^\flat}(\mathsf{x}_k)$ for some $b\in[0,1]$. Note that we can do so by the very definition of $j^\flat$. Then, we must have

$$\begin{aligned}\sum_{j\le j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k) - 1 &\ge \sum_{j<j^\flat} \widetilde{A}_{j,i}\cdot p_j(\mathsf{x}_k) + A_{j^\flat,i}\cdot p_{j^\flat}(\mathsf{x}_k) - 1\\ &= (1+\beta) - (1-b)\widetilde{A}_{j^\flat,i}\cdot p_{j^\flat}(\mathsf{x}_k) + A_{j^\flat,i}\cdot p_{j^\flat} - 1\\ &\ge \beta + b\cdot A_{j^\flat,i}\cdot p_{j^\flat}(\mathsf{x}_k)\ .\end{aligned}$$

Therefore, we conclude that

$$\begin{aligned}I &\ge \frac{\delta^\flat}{3}\Big(\sum_{j\le j^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k) - 1\Big) > \frac{\delta^\flat}{3}\cdot b\cdot A_{j^\flat,i}\cdot p_{j^\flat}(\mathsf{x}_k) = \frac{\mu\beta}{6\widetilde{A}_{j^\flat,i}}\cdot b\cdot \widetilde{A}_{j^\flat,i}\cdot p_{j^\flat}(\mathsf{x}_k)\\ &= \frac{\mu\beta}{6\widetilde{A}_{j^\flat,i}}\cdot\Big(\sum_{j\le j^\flat}\widetilde{A}_{j,i}\cdot p_j(\mathsf{x}_k) - (1+\beta)\Big) \ge \frac{\mu\beta}{12}\cdot u_i^*\cdot\Big(\sum_{j\le j^\flat}\widetilde{A}_{j,i}\cdot p_j(\mathsf{x}_k) - (1+\beta)\Big)\ .\end{aligned}$$

Above, the last inequality is because $u_i^*\cdot\widetilde{A}_{j^\flat,i} \le \langle \widetilde{A}_{j^\flat,:}, u^*\rangle \le 2$ by our definition of the adjusted $\widetilde{A}$.

Part *I′*. To lower bound I', consider every $j > j^\flat$ and the integral

$$\int_{\tau=0}^{\delta^\flat} A_{j,i}\cdot p_j(\mathsf{x}_k+\tau\mathbf{e}_i)d\tau\ .$$

Note that whenever $\tau \le \frac{\mu\beta}{2A_{j,i}} \le \frac{\mu\beta}{2A_{j^\flat,i}} = \delta^\flat$, we have that $p_j(\mathsf{x}_k+\tau\mathbf{e}_i) \ge p_j(\mathsf{x}_k)\cdot e^{-\beta/2} \ge \frac{1}{2}p_j(\mathsf{x}_k)$. Therefore, the above integral is at least $\frac{\mu\beta}{2A_{j,i}}\cdot A_{j,i}\cdot\frac{1}{2}p_j(\mathsf{x}_k)$. This implies a lower bound on I':

$$I' \ge \sum_{j>j^\flat}\frac{\mu\beta}{4A_{j,i}}\cdot A_{j,i}\cdot p_j(\mathsf{x}_k) \ge \frac{\mu\beta}{8}\cdot\sum_{j>j^\flat} u_i^*\cdot\widetilde{A}_{j,i}\cdot p_j(\mathsf{x}_k)\ ,$$

where again in the last inequality we have used $u_i^*\cdot\widetilde{A}_{j^\flat,i} \le \langle\widetilde{A}_{j^\flat,:}, u^*\rangle \le 2$ by our definition of $\widetilde{A}$.

Together. Combining the lower bounds on I and I', we obtain

$$f_\mu(\mathsf{x}_k) - f_\mu(\mathsf{y}_k^{(i)}) \ge I + I' \ge \frac{\mu\beta}{12}\cdot u_i^*\cdot\Big(\sum_{j=1}^m \widetilde{A}_{j,i}\cdot p_j(\mathsf{x}_k) - (1+\beta)\Big) = \frac{\mu\beta}{12}\cdot\langle -\widetilde{\eta}_k^{(i)}, u^*\rangle\ . \qquad \square$$

F Efficient Implementation of `PacLPSolver`

In this section, we illustrate how to implement each iteration of `PacLPSolver` to run in an expected $O(N/n)$ time. We maintain the following quantities

$$\mathsf{z}_k\in\mathbb{R}^n_{\ge 0},\quad \mathsf{az}_k\in\mathbb{R}^m_{\ge 0},\quad \mathsf{y}'_k\in\mathbb{R}^n,\quad \mathsf{ay}_k\in\mathbb{R}^m,\quad B_{k,1},B_{k,2}\in\mathbb{R}_+$$

throughout the algorithm, so as to ensure the following invariants are always satisfied

$$A\mathsf{z}_k = \mathsf{az}_k\ , \tag{F.1}$$

$$\mathsf{y}_k = B_{k,1} \cdot \mathsf{z}_k + B_{k,2} \cdot \mathsf{y}'_k \ , \qquad A\mathsf{y}_k = B_{k,1} \cdot A\mathsf{z}_k + B_{k,2} \cdot \mathsf{a}\mathsf{y}_k \ . \tag{F.2}$$

It is clear that when $k = 0$, letting $\mathsf{a}\mathsf{z}_k = A\mathsf{z}_0$, $\mathsf{y}'_k = \mathsf{y}_0$, $\mathsf{a}\mathsf{y}_k = A\mathsf{y}_0$, $B_{k,1} = 0$, and $B_{k,2} = 1$, we can ensure that all the invariants are satisfied initially. We denote $\|A_{:i}\|_0$ the number of nonzeros elements in vector $A_{:i}$. In each iteration $k = 1, 2, \ldots, T$:

- The step $\mathsf{x}_k = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$ does not need to be implemented.
- The value $\nabla_i f(\mathsf{x}_k)$ requires the knowledge of $p_j(\mathsf{x}_k) = e^{\frac{1}{\mu}((A\mathsf{x}_k)_j - 1)}$ for each j such that $A_{ij} \neq 0$. Accordingly, we need to know the value

  $$(A\mathsf{x}_k)_j = \tau(A\mathsf{z}_{k-1})_j + (1-\tau)(A\mathsf{y}_{k-1})_j = \big(\tau + (1-\tau)B_{k-1,1}\big)(A\mathsf{z}_{k-1})_j + (1-\tau)B_{k-1,2}\mathsf{a}\mathsf{y}_{k-1,j}$$

  for each such j. This can be computed in $O(1)$ time for each j, and $O(\|A_{:i}\|_0)$ time in total.
- Recall that the step $\mathsf{z}_k \leftarrow \arg\min_{z \in \Delta_{\mathsf{box}}} \left\{ \frac{1}{2}\|z - \mathsf{z}_{k-1}\|_A^2 + \langle n\alpha_k \xi_k^{(i)}, z \rangle \right\}$ can be written as $\mathsf{z}_k = \mathsf{z}_{k-1} + \delta \mathbf{e}_i$ for some $\delta \in \mathbb{R}$ that can be computed in $O(1)$ time (see Proposition 3.7). Observe also that $\mathsf{z}_k = \mathsf{z}_{k-1} + \delta\mathbf{e}_i$ yields $\mathsf{y}_k = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1} + \frac{\delta \mathbf{e}_i}{n\alpha_k L}$ due to Line 6 and Line 10 of Algorithm 1. Therefore, we can perform two explicit updates on z_k and $\mathsf{a}\mathsf{z}_k$ as

  $$\mathsf{z}_k \leftarrow \mathsf{z}_{k-1} + \delta \mathbf{e}_i \ , \quad \mathsf{a}\mathsf{z}_k \leftarrow A\mathsf{z}_{k-1} + \delta A_{:i}$$

  and two implicit updates on y_k as

  $$\begin{array}{lll} B_{k,1} = \tau + (1-\tau)B_{k-1,1} & , & B_{k,2} = (1-\tau)B_{k-1,2} \ , \\ \mathsf{y}'_k \leftarrow \mathsf{y}'_{k-1} + \delta\mathbf{e}_i \cdot \Big(-\frac{B_{k,1}}{B_{k,2}} + \frac{1}{n\alpha_k L}\frac{1}{B_{k,2}} \Big) & , & \mathsf{a}\mathsf{y}_k \leftarrow \mathsf{a}\mathsf{y}_{k-1} + \delta A_{:i} \cdot \Big(-\frac{B_{k,1}}{B_{k,2}} + \frac{1}{n\alpha_k L}\frac{1}{B_{k,2}} \Big) \end{array}$$

  It is not hard to verify that after these updates, we have

  $$\begin{aligned} \mathsf{y}_k &= B_{k,1} \cdot \mathsf{z}_k + B_{k,2} \cdot \mathsf{y}'_k = B_{k,1} \cdot \big(\mathsf{z}_{k-1} + \delta\mathbf{e}_i\big) + B_{k,2} \cdot \Big(\mathsf{y}'_{k-1} + \delta\mathbf{e}_i \cdot \Big(-\frac{B_{k,1}}{B_{k,2}} + \frac{1}{n\alpha_k L}\frac{1}{B_{k,2}}\Big)\Big) \\ &= B_{k,1} \cdot \mathsf{z}_{k-1} + B_{k,2} \cdot \Big(\mathsf{y}'_{k-1} + \delta\mathbf{e}_i \cdot \Big(\frac{1}{n\alpha_k L}\frac{1}{B_{k,2}}\Big)\Big) \\ &= B_{k,1} \cdot \mathsf{z}_{k-1} + B_{k,2} \cdot \mathsf{y}'_{k-1} + \frac{\delta\mathbf{e}_i}{n\alpha_k L} \\ &= \big(\tau + (1-\tau)B_{k-1,1}\big) \cdot \mathsf{z}_{k-1} + \big((1-\tau)B_{k-1,2}\big) \cdot \mathsf{y}'_{k-1} + \frac{\delta\mathbf{e}_i}{n\alpha_k L} \\ &= \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1} + \frac{\delta\mathbf{e}_i}{n\alpha_k L} \ . \end{aligned}$$

  One can similarly verify that $A\mathsf{y}_k = B_{k,1} \cdot A\mathsf{z}_k + B_{k,2} \cdot \mathsf{a}\mathsf{y}_k$ equals $A\mathsf{y}_k = \tau A\mathsf{z}_{k-1} + (1-\tau)A\mathsf{y}_{k-1} + \frac{\delta A e_i}{n\alpha_k L}$. In sum, these updates are dominated by the updates on $A\mathsf{z}_k$ and $\mathsf{a}\mathsf{y}_k$, each costing an $O(\|A_{:i}\|_0)$ running time, and ensure that the invariants in (F.1) and (F.2) are satisfied at iteration k.

In sum, we only need $O(\|A_{:i}\|_0)$ time to perform the updates in `PacLPSolver` for an iteration k if the coordinate i is selected. Therefore, each iteration of `PacLPSolver` can be implemented to run in an expected $O(\mathbf{E}_i[\|A_{:i}\|_0]) = O(N/n)$ time.

G Efficient Implementation of `CovLPSolver`

In this section we illustrate how to implement each iteration of `CovLPSolver` to run in an expected $O(N/n)$ time. We maintain the following quantities

$$\mathsf{z}'_k \in \mathbb{R}^n_+, \quad \mathsf{sz}_k \in \mathbb{R}_+, \quad \mathsf{sumz}_k \in \mathbb{R}_+, \quad \mathsf{az}_k \in \mathbb{R}^m_{\geq 0}, \quad \mathsf{y}'_k \in \mathbb{R}^n, \quad \mathsf{ay}_k \in \mathbb{R}^m, \quad B_{k,1}, B_{k,2} \in \mathbb{R}_+$$

throughout the algorithm, so as to maintain the following invariants are always satisfies

$$\mathsf{z}_k = \mathsf{z}'_k/\mathsf{sz}_k, \qquad \mathsf{sumz}_k = \mathbb{1}^T\mathsf{z}'_k, \qquad A\mathsf{z}_k = \mathsf{az}_k/\mathsf{sz}_k, \tag{G.1}$$

$$\mathsf{y}_k = B_{k,1}\cdot\mathsf{z}'_k + B_{k,2}\cdot\mathsf{y}'_k, \qquad A\mathsf{y}_k = B_{k,1}\cdot\mathsf{az}_k + B_{k,2}\cdot\mathsf{ay}_k \ . \tag{G.2}$$

It is clear that when $k=0$, letting $\mathsf{z}'_k = \mathsf{z}_0$, $\mathsf{sz}_k = 1$, $\mathsf{sumz}_k = \mathbb{1}^T\mathsf{z}_0$, $\mathsf{az}_k = A\mathsf{z}_0$, $\mathsf{y}'_k = \mathsf{y}_0$, $\mathsf{ay}_k = A\mathsf{y}_0$, $B_{k,1} = 0$, and $B_{k,2} = 1$, we can ensure that all the invariants are satisfied initially.

We denote by $\|A_{:i}\|_0$ the number of nonzero elements in vector $A_{:i}$. In each iteration $k = 1, 2, \ldots, T$:

- The step $\mathsf{x}_k = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1}$ does not need to be implemented.
- The value $p_j(\mathsf{x}_k) = e^{\frac{1}{\mu}(1-(A\mathsf{x}_k)_j)}$ for each j only requires the knowledge of

  $$(A\mathsf{x}_k)_j = \tau(A\mathsf{z}_{k-1})_j + (1-\tau)(A\mathsf{y}_{k-1})_j = \big(\tau + (1-\tau)B_{k-1,1}\big)\frac{\mathsf{az}_{k-1,j}}{\mathsf{sz}_{k-1}} + (1-\tau)B_{k-1,2}\mathsf{ay}_{k-1,j} \ .$$

  This can be computed in $O(1)$ time.
- The value $\nabla_i f(\mathsf{x}_k)$ requires the knowledge of $p_j(\mathsf{x}_k)$ for each $j \in [m]$ such that $A_{ij} \neq 0$. Since we have $\|A_{:i}\|_0$ such j's, we can compute $\nabla_i f(\mathsf{x}_k)$ in $O(\|A_{:i}\|_0)$ time.
- Letting $\delta = (1+\gamma)n\alpha_k\xi^{(i)}_{k,i}$, recall that the mirror step $\mathsf{z}_k \leftarrow \arg\min_{z\in\Delta_{\mathsf{simplex}}}\big\{V_{\mathsf{z}_{k-1}}(z) + \langle\delta\mathbf{e}_i, z\rangle\big\}$ has a very simple form (see Proposition 6.12): first multiply the i-th coordinate of z_{k-1} by $e^{-\delta}$ and then, if the sum of all coordinates have exceeded $2\mathsf{OPT}'$, scale everything down so as to sum up to $2\mathsf{OPT}'$. This can be implemented as follows: setting $\delta_1 = \mathsf{z}'_{k-1,i}(e^{-\delta}-1)$,

  $$\begin{aligned} &\mathsf{z}'_k \leftarrow \mathsf{z}'_{k-1} + \delta_1\mathbf{e}_i \quad , \quad \mathsf{az}_k \leftarrow \mathsf{az}_{k-1} + \delta_1 A_{:i} \ , \\ &\mathsf{sumz}_k \leftarrow \mathsf{sumz}_{k-1} + \delta_1 \quad , \quad \mathsf{sz}_k \leftarrow \mathsf{sz}_k \cdot \max\Big\{1, \tfrac{\mathsf{sumz}_k}{\mathsf{sz}_{k-1}\cdot 2\mathsf{OPT}'}\Big\} \ . \end{aligned}$$

  These updates can be implemented to run in $O(\|A_{:i}\|_0)$ time, and they together ensure that the invariants in (G.1) are satisfied at iteration k.
- Recall that the gradient step is of the form $\mathsf{y}_k \leftarrow \mathsf{x}_k + \delta_2\cdot\mathbf{e}_i$ for some value $\delta_2 \geq 0$. This value δ_2 can be computed in $O(\|A_{:i}\|_0)$ time, since each $p_j(\mathsf{x}_k)$ can be computed in $O(1)$ time, and we can sort the rows of each column of A by preprocessing.

  Since $\mathsf{y}_k = \mathsf{x}_k + \delta_2\cdot\mathbf{e}_i = \tau\mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1} + \delta_2\mathbf{e}_i$, we can implement this update by letting

  $$\begin{aligned} &B_{k,1} = \tfrac{\tau}{\mathsf{sz}_{k-1}} + (1-\tau)B_{k-1,1} \quad , \quad B_{k,2} = (1-\tau)B_{k-1,2} \\ &\mathsf{y}'_k \leftarrow \mathsf{y}'_{k-1} + \mathbf{e}_i\cdot\Big(-\tfrac{B_{k,1}\delta_1}{B_{k,2}} + \tfrac{\delta_2}{B_{k,2}}\Big) \quad , \quad \mathsf{ay}_k \leftarrow \mathsf{ay}_{k-1} + A_{:i}\cdot\Big(-\tfrac{B_{k,1}\delta_1}{B_{k,2}} + \tfrac{\delta_2}{B_{k,2}}\Big) \end{aligned}$$

  It is not hard to verify that after these updates, we have

  $$\mathsf{y}_k = B_{k,1}\cdot\mathsf{z}'_k + B_{k,2}\cdot\mathsf{y}'_k = B_{k,1}\cdot\big(\mathsf{z}'_{k-1} + \delta_1\mathbf{e}_i\big) + B_{k,2}\cdot\Big(\mathsf{y}'_{k-1} + \mathbf{e}_i\cdot\Big(-\frac{B_{k,1}\delta_1}{B_{k,2}} + \frac{\delta_2}{B_{k,2}}\Big)\Big)$$

$$
\begin{aligned}
&= B_{k,1} \cdot \mathsf{z}'_{k-1} + B_{k,2} \cdot \big(\mathsf{y}'_{k-1} + \delta_2 \mathbf{e}_i / B_{k,2}\big) \\
&= B_{k,1} \cdot \mathsf{z}'_{k-1} + B_{k,2} \cdot \mathsf{y}'_{k-1} + \delta_2 \mathbf{e}_i \\
&= \big(\frac{\tau}{\mathsf{sz}_{k-1}} + (1-\tau) B_{k-1,1}\big) \cdot \mathsf{z}'_{k-1} + \big((1-\tau) B_{k-1,2}\big) \cdot \mathsf{y}'_{k-1} + \delta_2 \mathbf{e}_i \\
&= \tau \mathsf{z}_{k-1} + (1-\tau)\mathsf{y}_{k-1} + \delta_2 \mathbf{e}_i \ .
\end{aligned}
$$

One can similarly verify that $A\mathsf{y}_k = B_{k,1} \cdot \mathsf{az}_k + B_{k,2} \cdot \mathsf{ay}_k$ equals $A\mathsf{y}_k = \tau A\mathsf{z}_{k-1} + (1-\tau)A\mathsf{y}_{k-1} + \delta_2 A_{:i}$. These updates can be implemented to run in $O(\|A_{:i}\|_0)$ time, and they together ensure that the invariants in (G.2) are satisfied at iteration k.

In sum, we only need $O(\|A_{:i}\|_0)$ time to perform the updates in `CovLPSolver` for an iteration k if the coordinate i is selected. Therefore, each iteration of `CovLPSolver` can be implemented to run in an expected $O(\mathbf{E}_i[\|A_{:i}\|_0]) = O(N/n)$ time.

References

[1] Zeyuan Allen-Zhu, Zhenyu Liao, and Lorenzo Orecchia. Spectral Sparsification and Regret Minimization Beyond Multiplicative Updates. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, STOC '15, 2015.

[2] Zeyuan Allen-Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. *ArXiv e-prints*, abs/1407.1537, July 2014.

[3] Zeyuan Allen-Zhu and Lorenzo Orecchia. Nearly-Linear Time Positive LP Solver with Faster Convergence Rate. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, STOC '15, 2015.

[4] Zeyuan Allen-Zhu and Lorenzo Orecchia. Using optimization to break the epsilon barrier: A faster and simpler width-independent algorithm for solving positive linear programs in parallel. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, 2015.

[5] Zeyuan Allen-Zhu, Peter Richtárik, Zheng Qu, and Yang Yuan. Even faster accelerated coordinate descent using non-uniform sampling. In *ICML*, 2016.

[6] Sanjeev Arora, Elad Hazan, and Satyen Kale. The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory of Computing*, 8:121–164, 2012.

[7] Baruch Awerbuch and Rohit Khandekar. Stateless distributed gradient descent for positive linear programs. *Proceedings of the fourtieth annual ACM symposium on Theory of computing - STOC 08*, page 691, 2008.

[8] Baruch Awerbuch, Rohit Khandekar, and Satish Rao. Distributed algorithms for multicommodity flow problems via approximate steepest descent framework. *ACM Transactions on Algorithms*, 9(1):1–14, December 2012.

[9] Yair Bartal, John W. Byers, and Danny Raz. Global optimization using local information with applications to flow control. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 303–312. IEEE Comput. Soc, 1997.

[10] Yair Bartal, John W. Byers, and Danny Raz. Fast, Distributed Approximation Algorithms for Positive Linear Programming with Applications to Flow Control. *SIAM Journal on Computing*, 33(6):1261–1279, January 2004.

[11] D. Bienstock and G. Iyengar. Faster approximation algorithms for packing and covering problems. Technical report, Columbia University, September 2004. Preliminary version published in STOC '04.

[12] John Byers and Gabriel Nasser. Utility-based decision-making in wireless sensor networks. In *Mobile and Ad Hoc Networking and Computing, 2000. MobiHOC. 2000 First Annual Workshop on*, pages 143–144. IEEE, 2000.

[13] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, 2006.

[14] Fabián A. Chudak and Vânia Eleutério. Improved Approximation Schemes for Linear Programming Relaxations of Combinatorial Optimization Problems. In *Proceedings of the 11th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 81–96, 2005.

[15] Ran Duan and Seth Pettie. Linear-Time Approximation for Maximum Weight Matching. *Journal of the ACM*, 61(1):1–23, January 2014.

[16] Olivier Fercoq and Peter Richtárik. Accelerated, Parallel and Proximal Coordinate Descent. *ArXiv e-prints*, abs/1312.5799:25, December 2013.

[17] Lisa K. Fleischer. Approximating Fractional Multicommodity Flow Independent of the Number of Commodities. *SIAM Journal on Discrete Mathematics*, 13(4):505–520, January 2000.

[18] Naveen Garg and Jochen Könemann. Faster and Simpler Algorithms for Multicommodity Flow and Other Fractional Packing Problems. *SIAM Journal on Computing*, 37(2):630–652, January 2007.

[19] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Journal of the ACM (JACM)*, 58(6):30, 2011.

[20] Philip Klein and Neal Young. On the number of iterations for dantzig-wolfe optimization and packing-covering approximation algorithms. In Gérard Cornuéjols, Rainer E. Burkard, and Gerhard J. Woeginger, editors, *Integer Programming and Combinatorial Optimization*, volume 1610 of *Lecture Notes in Computer Science*, pages 320–327. Springer Berlin Heidelberg, 1999.

[21] Christos Koufogiannakis and Neal E. Young. A Nearly Linear-Time PTAS for Explicit Fractional Packing and Covering Linear Programs. *Algorithmica*, pages 494–506, March 2013. Previously appeared in FOCS '07.

[22] Michael Luby and Noam Nisan. A parallel approximation algorithm for positive linear programming. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing - STOC '93*, pages 448–457, New York, New York, USA, 1993. ACM Press.

[23] Aleksander Madry. Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms. In *Proceedings of the 42nd ACM symposium on Theory of computing - STOC '10*, page 121, New York, New York, USA, 2010. ACM Press.

[24] Arkadi Nemirovski. Prox-Method with Rate of Convergence $O(1/t)$ for Variational Inequalities with Lipschitz Continuous Monotone Operators and Smooth Convex-Concave Saddle Point Problems. *SIAM Journal on Optimization*, 15(1):229–251, January 2004.

[25] Yu Nesterov. Rounding of convex sets and efficient gradient methods for linear programming problems. *Optimisation Methods and Software*, 23(1):109–128, 2008.

[26] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. In *Doklady AN SSSR (translated as Soviet Mathematics Doklady)*, volume 269, pages 543–547, 1983.

[27] Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, December 2005.

[28] Yurii Nesterov. Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems. *SIAM Journal on Optimization*, 22(2):341–362, jan 2012.

[29] Serge A. Plotkin, David B. Shmoys, and Éva Tardos. Fast Approximation Algorithms for Fractional Packing and Covering Problems. *Mathematics of Operations Research*, 20(2):257–301, May 1995.

[30] Luca Trevisan. Parallel Approximation Algorithms by Positive Linear Programming. *Algorithmica*, 21(1):72–88, May 1998.

[31] Di Wang, Michael W. Mahoney, Nishanth Mohan, and Satish Rao. Faster parallel solver for positive linear programs via dynamically-bucketed selective coordinate descent. *ArXiv e-prints*, abs/1511.06468, November 2015.

[32] Di Wang, Satish Rao, and Michael W. Mahoney. Unified acceleration method for packing and covering problems via diameter reduction. In *ICALP*, 2016.

[33] Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In *42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS'01)*, pages 538–546. IEEE Comput. Soc, 2001.

[34] Neal E. Young. Nearly linear-time approximation schemes for mixed packing/covering and facility-location linear programs. *ArXiv e-prints*, abs/1407.3015, July 2014.

[35] Edo Zurel and Noam Nisan. An efficient approximate allocation algorithm for combinatorial auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 125–136. ACM, 2001.